

INTERPRETING THE UNRESOLVED INTENSITY OF COSMOLOGICALLY REDSHIFTED LINE RADIATION

E. R. SWITZER¹, T.-C. CHANG², K. W. MASUI^{3,4}, U.-L. PEN^{5,4}, T. C. VOYTEK^{6,7}*Submitted to ApJ*

ABSTRACT

Intensity mapping experiments survey the spectrum of diffuse line radiation rather than detect individual objects at high signal-to-noise ratio. Spectral maps of unresolved atomic and molecular line radiation contain three-dimensional information about the density and environments of emitting gas and efficiently probe cosmological volumes out to high redshift. Intensity mapping survey volumes also contain all other sources of radiation at the frequencies of interest. Continuum foregrounds are typically $\sim 10^2$ – 10^3 times brighter than the cosmological signal. The instrumental response to bright foregrounds will produce new spectral degrees of freedom that are not known in advance, nor necessarily spectrally smooth. The intrinsic spectra of foregrounds may also not be well known in advance. We describe a general class of quadratic estimators to analyze data from single-dish intensity mapping experiments and determine contaminated spectral modes from the data themselves. The key attribute of foregrounds is not that they are spectrally smooth, but instead that they have fewer bright spectral degrees of freedom than the cosmological signal. Spurious correlations between the signal and foregrounds produce additional bias. Compensation for signal attenuation must estimate and correct this bias. A successful intensity mapping experiment will control instrumental systematics that spread variance into new modes, and it must observe a large enough volume that contaminant modes can be determined independently from the signal on scales of interest.

Subject headings: methods: data analysis – methods: statistical – (cosmology:) diffuse radiation – (cosmology:) large-scale structure of universe

1. INTRODUCTION

Intensity mapping is an emerging technique for cosmological observation. It uses atomic or molecular transition radiation to tomographically map large volumes of the universe. These data volumes contain a combination of information about the abundance, environment, and velocity of emitters. Tomographic line surveys offer the potential to capture the dynamic universe in epochs that are otherwise difficult to observe (Madau et al. 1997) and to probe a variety of galactic environments in aggregate (Uzgil et al. 2014; Li et al. 2015). The recovery of many modes could allow intensity surveys to compete with dark energy constraints from standard spectroscopic galaxy surveys (Chang et al. 2008; Loeb & Wyithe 2008).

Intensity mapping shares some parallels with studies of cosmic background radiation and spectroscopic galaxy surveys. Like background radiation studies, intensity mapping experiments use sensitive receivers to map diffuse emission. In contrast, a survey to resolve the individual sources of emission requires significantly higher sensitivity and resolution, both of which drive costs or reduce scope. Intensity mapping

only requires resolution to reach cosmologically interesting scales, or scales with large enough fluctuations to secure a detection. Additionally, an intensity mapping survey is sensitive to the integral of the luminosity function, taking advantage of all the emitted radiation that is available, not just the brightest sources. However, the lack of source discrimination also makes intensity mapping survey volumes significantly more difficult to interpret than spectroscopic galaxy surveys. An intensity survey generally has not just the line emission but also all other sources of continuum and transition radiation from other redshifts.

This paper describes a method for estimating the power spectrum of intensity mapping volumes, subject to bright foreground emission and instrumental response. We inherit from the framework (Tegmark 1997; Tegmark et al. 1998) of optimal quadratic estimators and extend methods that can be applied to single-dish intensity surveys (Liu & Tegmark 2011; Dillon et al. 2013, 2014; Switzer et al. 2013). Rather than using a fully optimal estimator, we construct a more generic estimator and new methods for handling the impact of the instrumental beam and foreground cleaning. Treatment of foregrounds is our primary focus, which in a nutshell translates into specifying the most effective foreground covariance matrix to down-weight contamination. We argue that the data themselves are the best source of information about foreground covariance, especially in light of the instrument's response to bright foregrounds.

Intensity mapping was originally developed for 21 cm radiation (Hogan & Rees 1979; Scott & Rees 1990; Madau et al. 1997), but has been studied for several other lines (in increasing frequency): deuterium (Sigurdson & Furlanetto 2006), $^3\text{He}+$ (McQuinn & Switzer 2009; Bagla & Loeb 2009), CO (Righi et al. 2008; Carilli 2011; Lidz et al. 2011; Breyse et al. 2014; Li et al. 2015), C II (Gong et al. 2012; Silva et al. 2014; Uzgil et al. 2014; Yue et al. 2015), Ly α

¹ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA² Academia Sinica Institute of Astronomy and Astrophysics 11F of Astro-Math Building, AS/NTU, 1 Roosevelt Rd Sec 4, Taipei, 10617, Taiwan³ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Rd. Vancouver, V6T 1Z1, Canada⁴ Canadian Institute for Advanced Research, CIFAR Program in Cosmology and Gravity, Toronto, ON, M5G 1Z8⁵ Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada⁶ McWilliams Center for Cosmology, Carnegie Mellon University, Department of Physics, 5000 Forbes Ave., Pittsburgh, PA 15213, USA⁷ Astrophysics and Cosmology Research Unit, School of Chemistry and Physics, University of KwaZulu-Natal, Durban, 4041, South Africa
eric.r.switzer@nasa.gov

(Gong et al. 2014; Pullen et al. 2014), C and O fine-structure (Kusakabe & Kawasaki 2012), and X-ray lines (Hütsi et al. 2012).

Many experiments to search for redshifted line emission are planned or under way. These experiments deploy a wide range of technologies to observe across the frequency range of the lines of interest, from the present to the dark ages. We will focus specifically on non-interferometric, “single-dish” methods that have a common set of simpler instrumental considerations. “Single dish” will refer to any single-aperture optical path, including refractive designs. Throughout, the beam or point-spread functions are constant in time, axisymmetric, frequency dependent and may also have some off-diagonal Mueller mixing. While there are many parallels in interferometers, they are beyond the scope of this work (see, e.g. Dillon et al. (2014); Shaw et al. (2014a,b); Ali et al. (2015)).

Aside from the 21 cm transition, most proposed or active intensity mapping experiments use a single-dish architecture. CO has been sought by Pullen et al. (2013) and proposed by COMAP (Li et al. 2015). TIME (Crites et al. 2014) and SPHEREX (Doré et al. 2014) are proposed for C II and Ly α respectively. Within 21 cm efforts, single-dish instruments have been used (GBT, Chang et al. (2010); Masui et al. (2013); Switzer et al. (2013)) or proposed (BINGO, Battye et al. (2012)) for studies at $z \sim 1$. The methods described here were developed for GBT studies. Even at low redshifts, 21 cm interferometers such as BAOBAB (Pober et al. 2013) and CHIME (Bandura et al. 2014) are needed to compete with dark energy constraints from optical galaxy surveys. Interferometers are the only realistic methodology for 21 cm studies of reionization.

Diffuse radiation from cosmologically redshifted atomic transitions has only been conclusively detected in cross-correlation with spectroscopic galaxy surveys. The cross-power with a density field is not biased by foregrounds, which instead boost errors (assuming that foregrounds are unrelated to signal). Croft et al. (2015) recently detected Ly α emission intensity in cross-correlation with BOSS quasars from $z = 2 - 3.5$. Masui et al. (2013) detected 21 cm radiation at $z \sim 1$ in cross-correlation between dedicated GBT observations and the WiggleZ survey (Drinkwater et al. 2010) and inferred the 21 cm contribution to the auto-power (Switzer et al. 2013). Bounds on the auto-power at modest redshift date to Bebbington (1986). In the absence of a coeval spectroscopic galaxy survey, such as at reionization, cross-correlation with intensity maps of other atomic lines (e.g. Visbal & Loeb (2010)) could secure a detection of cosmological structure, up to challenges of correlated foregrounds. As we will argue below, the principal challenge of intensity mapping experiments is that line radiation only makes up $\sim 10^{-2} - 10^{-3}$ of the intensity of fluctuations in continuum radiation at most frequencies of interest.

The optimal estimator for the power spectrum requires the covariance matrix of the maps, and specification of this covariance is a central challenge of analysis of intensity mapping data. While the non-Gaussianity of foregrounds could be distinguishable from the near-Gaussianity of the signals, we will not consider separation using higher-point statistics. The full foreground covariance of a 3D survey is an $N_{\text{pix}} \times N_{\text{pix}}$ matrix for N_{pix} total map pixels and hence already requires an enormous amount of information about foregrounds and the instrument response to foregrounds. We have little prior knowledge of either.

Common approaches to specifying foreground covari-

ance amount to different forms of dimensionality reduction. Liu & Tegmark (2011) show that most of the covariance that distinguishes foregrounds is in the ν, ν' directions rather than combinations involving angular separations. Fitting polynomials along the lines of sight corresponds to an ansatz of ν, ν' -only covariance contributed by those polynomial modes (e.g. Wang et al. (2006)). These ansatzes can be better tuned to astrophysical foregrounds by using models of the emission. Astrophysical synchrotron intensity is thought to be described by a limited number of spectral modes (e.g. Liu & Tegmark (2012)).

If foregrounds are bright and the instrument response is not sufficiently well understood, then misspecification of the foreground covariance can result in significant contamination in the maps that is not down-weighted. For foregrounds 10^3 times the signal, an unattributed 1% error in calibration at one frequency could result in contamination that is 10 times larger than the signal. Worse still, if the spectral calibration varies in time, then each line of sight effectively sees a different bright foreground. These may even form a complete basis of bright spectral modes, making the signal indistinguishable. Faraday rotation through polarization leakage into intensity in radio surveys is another example (Moore et al. 2013) where additional spectral modes can be produced by the instrumental response.

The position we take in this paper is that (1) some best-effort calibration has been applied, but that there is residual structure in the map related to the instrument, and (2) the intrinsic foreground cannot be modeled well in advance. In the regime of high foregrounds, measurements of foregrounds from other instruments or wavelengths may not have the fidelity to be useful for foreground subtraction. Intensity mapping surveys will often be the deepest surveys available in a region. This is ultimately due to the dimness of the atomic radiation, but another factor is that many 2D surveys stop integrating beyond the confusion limit. A typical intensity mapping experiment can benefit from thermal noise levels well below spatial confusion. These factors argue that the intensity survey volumes will be the best sources of information about foreground covariance rather than prior models of the instrument or intrinsic foreground emission.

Even if the relevant foreground covariance is separable as ν, ν' blocks, we are unlikely to estimate that covariance matrix at full rank from independent sight lines in the data. The final dimensionality reduction we assume is that the ν, ν' covariance estimated from the data will have a few dominant eigenvectors that are measured with high signal-to-noise ratio, while the remaining data are dominated by the cosmological signal and thermal noise. The foreground eigenvectors are spectral degrees of freedom that can be projected out of each line of sight. Determination of contaminated modes in the data themselves has been exploited in GBT data (Chang et al. 2010; Masui et al. 2013; Switzer et al. 2013), GMRT (Paciga et al. 2013), and most recently in PAPER (e.g. Ali et al. (2015)). Blind methods have been considered for SKA (Wolz et al. 2014; Alonso et al. 2015) and BINGO (Bigot-Sazy et al. 2015). Switzer & Liu (2014) develop a similar method for monopole signals. The success of this method relies on (1) whether instrument response to bright foregrounds can be explained by fewer spectral modes than the cosmological signal and (2) whether these modes can be determined from the data themselves, independently from the signal. These requirements relate to the rank of foregrounds rather than spectral smoothness.

The first requirement is intuitive and says that the signal must have degrees of spectral variation that are orthogonal to the foreground modes that are removed. The second requirement arises from the fact that the foreground modes are determined from the data themselves, which have foreground, signal, and noise. Spurious correlations between foreground and signal result in residual covariance in the maps that is anticorrelated with the cosmological signal. This general effect is familiar from the ILC bias (see, e.g. Efstathiou et al. (2009)). Especially on scales that are large compared to the survey volume, there are too few realizations of signal fluctuations for the spurious correlation with the foregrounds to “average down.” In this case, it is difficult to disentangle signal from foreground.

Section 2 briefly reviews continuum foreground levels for several transitions of interest and describes data from the GBT-wide survey (Switzer et al. 2013) that we use as an example throughout. Section 3 develops the general (potentially suboptimal) quadratic estimator, the skeleton prescribed by the optimal estimator, and ways of calibrating the analysis using Monte Carlo simulations. Section 4 builds up the formalism of mode removal, line-of-sight cleaning, and rules of thumb for the impact of cleaning known modes. Rather than known modes, spectral contamination can be determined from the data themselves (Section 5), leading to modifications for rules of thumb of signal loss and estimation of transfer functions. Finally, Sections 6 and 7 describe the procedure for assembling final power spectral estimates using subseason cross-powers, weighted 2D to 1D averaging, and development of errors.

2. CONTINUUM FOREGROUNDS AND DATA USED

2.1. Review of Continuum Foreground Levels

This section briefly reviews the literature and estimates magnitudes of continuum foregrounds. The goal is not to make a precise determination, but instead to argue that intensity mapping has a common set of challenges. In detail, foreground challenges will vary across the bands and survey strategies. For example, an experiment that needs wide sky areas may not be able to avoid bright extragalactic sources, galactic emission, or zodiacal light, while small areas can be better tuned. To get a rough understanding of irreducible continuum levels, we will consider fluctuations in extragalactic radiation on angular scales where the signal has order-unity fluctuations. Here the mean line emission intensity serves as a proxy. Smaller angular scales may have higher signal variance, but a study of signal vs. continuum emission on different angular scales is deferred to future work for particular lines and redshifts.

Considering 21 cm first, Masui et al. (2013) show intensity maps at 800 MHz ($z \approx 0.8$) with foreground fluctuations of order a kelvin, while the signal fluctuations are ~ 0.2 mK. While the 21 cm reionization signal is brighter, the foregrounds are commensurately brighter because of the synchrotron spectral index, yielding a similar challenge. Considering CO(1-0) at 115 GHz, Breyse et al. (2014) and Li et al. (2015) suggest mean temperatures of ~ 1 μ K at $z \sim 3$. The scales of interest at tens of arcminutes are reasonably analogous to the Cosmic Background Imager (CBI) (Pearson et al. 2003), which finds fluctuations in the raw maps at the level of several hundred μ K at 30 GHz. (Note that some sources could be cleaned or masked based on catalogs, but cosmic microwave background [CMB] would remain.) By $z \sim 8$, extra-

galactic and galactic synchrotron become more problematic, while the mean brightness is expected to be a similar order of magnitude (Lidz et al. 2011).

Moving to C II (157.7 μ m) at reionization, Silva et al. (2014) estimate a mean intensity of 4×10^2 Jy sr $^{-1}$ for $z = 5.3 - 8.5$ (300–200 GHz), while the extragalactic fluctuations on scales of several arcmin in the Atacama Cosmology Telescope (ACT) (Dünner et al. 2013) and South Pole Telescope (SPT) (Schaffer et al. 2011) data in this band are ~ 100 μ K at 220 GHz, or $\sim 5 \times 10^4$ Jy sr $^{-1}$ (CMB dominated). C II emission is thought to reach a maximum at $z \approx 1$, at $\sim 5 \times 10^3$ Jy sr $^{-1}$ (Uzgil et al. 2014), while *Herschel* ATLAS (Eales et al. 2010) shows cosmic infrared background fluctuations ranging over $\sim 2 \times 10^6$ Jy sr $^{-1}$ at 350 μ m.

For Ly α (10.1 eV) mapping, Croft et al. (2015) measure the mean surface brightness in cross-correlation between quasars and spectra, finding $\nu I_\nu = (0.74 \pm 0.17)$ nWm $^{-2}$ sr $^{-1}$ (at 4500 Å) across $z = 2 - 3.5$, $\sim 21 - 35$ higher than previously expected (see, e.g. Pullen et al. (2014)). These redshifts span the expected peak of emission from high star formation rates. Reported mean backgrounds in the BOSS spectra include all sources of radiation, including terrestrial, and are $\nu I_\nu \sim 50$ nWm $^{-2}$ sr $^{-1}$. For lower redshifts, GALEX (Murthy et al. 2010) data suggest astrophysical backgrounds of ~ 10 nWm $^{-2}$ sr $^{-1}$ for 5.1–8.4 eV, similar to the general cosmic optical background (Hauser & Dwek 2001). In the reionization era where Ly α has shifted to $\sim \mu$ m, Levenson & Wright (2008) find that the extragalactic mean contribution at 3.6 μ m is 9 nWm $^{-2}$ sr $^{-1}$.

In summary, typical continuum contamination to intensity surveys is $10^2 - 10^3$ times the line contribution. The underlying challenge is the small fraction of total luminosity emitted through line radiation. Instrumental response needs to be controlled at a subpercent level, commensurate with the brightness of the foregrounds.

We will consider emission of a single line rather than a more general SED (de Putter et al. 2014) and neglect interlopers at other redshifts (e.g. Breyse et al. (2015)), which have received more attention than instrumental response to bright continua.

2.2. Data Used to Demonstrate the Method

We will use data from the GBT-wide survey to give context to the estimator described here. Previous publications review the observations and describe the power spectrum of GBT data (Switzer et al. 2013) and the cross-correlation (Masui et al. 2013) with the WiggleZ survey (Blake et al. 2011). No results from new data are reported. Masui et al. (2013) describe the observations in more detail. The GBT-wide intensity survey used the prime-focus receiver to map a $\sim 7^\circ \times 4.3^\circ$ region from 700 to 900 MHz with FWHM $\sim 0.3^\circ$ and 256 spectral bands.

Our starting point will be the map that has been estimated from time-ordered data. A framework for estimating maps is well established from CMB analysis and depends on particulars such as noise correlations and frequency masking. The details of the calibration, radio frequency interference (RFI) mitigation, and mapmaking used to produce the maps used here can be found in Masui (2013). Section 5.1 describes the calibration of GBT data in regard to the spectral structure of contaminated modes.

We use Gaussian signal realizations of the Empirical-NL model of Blake et al. (2011), which uses HALOFIT

(Smith et al. 2003) for nonlinear power, Kaiser redshift distortions, and streaming of the Lorentz form with $\sigma_v = 300 h \text{ km/s}$. To agree with Empirical-NL, we use $\Omega_m = 0.27$, $\Omega_\Lambda = 0.73$, $\Omega_b/\Omega_m = 0.166$, $h = 0.72$, and $n_s = 0.96$ from Komatsu et al. (2009). These parameters are also used to translate the observed regions into comoving Cartesian coordinates. We use approximations to the growth factor from Kasai (2010). The brightness temperature of the 21 cm line is taken to be

$$T_b(z) = T_o \frac{\Omega_{HI}(z)}{10^{-3}} \left[\frac{\Omega_m + \Omega_\Lambda(1+z)^{-3}}{0.29} \right]^{-1/2} \left[\frac{1+z}{2.5} \right]^{1/2}, \quad (1)$$

with $T_o = 0.39 \text{ mK}$.

3. THE QUADRATIC ESTIMATOR

3.1. Map Notations

The intensity survey produces maps at several frequencies. We can represent these maps as a matrix \mathbf{X} with dimensions $N_\nu \times N_\theta$, where N_ν is the number of frequency slices and N_θ is the number of angular pixels observed. This is a stack of 2D maps, where the map at each frequency is unraveled into an N_θ -long vector. An alternative is to unravel the entire 3D volume into a single vector \mathbf{x} of length $N_\nu \cdot N_\theta$. These representations are related through the operation $\text{vec}(\mathbf{X}) = \mathbf{x}$ which unravels the stacked map matrix into a vector. The matrix form \mathbf{X} has the useful property that operations can act explicitly on either the frequency or angular side as $\mathbf{A}\mathbf{X}\mathbf{B}$, for \mathbf{A} and \mathbf{B} here, respectively. While this looks like a quadratic conjugation of \mathbf{X} , it is linear in \mathbf{x} through the relation

$$\text{vec}(\mathbf{A}\mathbf{X}\mathbf{B}) = (\mathbf{B}^T \otimes \mathbf{A})\mathbf{x}, \quad (2)$$

where \otimes is the Kronecker product. Another perspective is that if a matrix \mathbf{W} multiplying \mathbf{x} can be written separably as $\mathbf{W}_\nu \otimes \mathbf{W}_\theta$, then those weights can act independently on the spectral and spatial axes of \mathbf{X} . We will use the \mathbf{X} representation of the data when it is convenient to call out this separable form, and \mathbf{x} when we want a simple “cascaded” set of operations to apply to the full map. Note that 2D intensity surveys directly analogous to CMB cross-correlation can also be pursued (see, e.g. Pullen et al. (2013)); however, all cases studied here will be 3D.

3.2. Quadratic Estimators

In this section we review the general quadratic estimator, following early work (Tegmark 1997) and recent applications to intensity mapping (Liu & Tegmark 2011, 2012; Dillon et al. 2013, 2014; Switzer et al. 2013). We break with tradition somewhat by deriving expressions for a general quadratic estimator rather than specializing to the optimal case.

Take a covariance model with Gaussian thermal noise \mathbf{N} (initially ignoring foregrounds) and signal that is decomposed as a set of amplitudes p_α times the covariance $\mathbf{C}_{,\alpha}$ of modes in the map. Throughout, commas will denote derivatives. Then

$$\mathbf{C} = \langle \mathbf{x}\mathbf{x}^T \rangle = \mathbf{N} + \sum_\alpha p_\alpha \mathbf{C}_{,\alpha}. \quad (3)$$

We will develop a particular $\mathbf{C}_{,\alpha}$ for (k_\perp, k_\parallel) modes of the power spectra in Section 3.3.

Our goal is to estimate the amplitudes p_α of the covariance from a given map \mathbf{x} , and infer their errors and correlations.

A general class of covariance estimators forms a quadratic combination of the data, subtracts a bias b_α and then takes a linear combination of the $\hat{q}_\alpha = \hat{q}|_\alpha$, as

$$\begin{aligned} \hat{q}_\alpha &= \mathbf{x}^T \mathbf{Q}_\alpha \mathbf{x} - b_\alpha \\ \hat{\mathbf{p}} &= \mathbf{R} \hat{\mathbf{q}}. \end{aligned} \quad (4)$$

The expectation value of the quadratic combination is

$$\langle \mathbf{x}^T \mathbf{Q}_\alpha \mathbf{x} \rangle = \text{Tr}(\mathbf{C} \mathbf{Q}_\alpha) \quad (5)$$

and is sensitive to the variance of the noise \mathbf{N} as well as the signal $\sum_\alpha p_\alpha \mathbf{C}_{,\alpha}$. By choosing to subtract a noise bias $b_\alpha = \text{Tr}(\mathbf{N} \mathbf{Q}_\alpha)$ based on a model for \mathbf{N} , \hat{q}_α measures just the signal covariance. Finally, the matrix \mathbf{R} takes a linear combination of the band powers \hat{q}_α to form a final estimate \hat{p}_α . The vector $\hat{\mathbf{q}}$ contains pseudo-powers in the language of Hivon et al. (2002). While it has mainly been described in a role of decorrelating (Hamilton & Tegmark 2000) band powers, the final linear combination of band powers by \mathbf{R} performs several roles as follows: (1) a normalization to ensure that \mathbf{Q}_α recovers an unbiased estimate of the signal (neglecting beam and foreground considerations), (2) a correction for signal attenuation from beam convolution and foreground down-weighting, and (3) an operation that decorrelates band powers.

In the standard treatment of quadratic estimators (see, e.g. Tegmark (1997)), one seeks to minimize the variance of the estimator $\hat{\mathbf{p}}$ subject to the Lagrange constraint that it is an unbiased estimate of true \mathbf{p} . Most literature develops expressions henceforth assuming an optimal estimator. In the case of the optimal estimator, the Fisher matrix is ubiquitous because the estimator can saturate the Cramer-Rao bound for Gaussian fields. In contrast, we will assume that \mathbf{Q}_α is given and will in general be suboptimal. Section 3.3 derives the form of the optimal estimator to develop some intuition for good suboptimal estimators. The fully optimal estimator requires a complete model of the covariance for its optimal weights. In any near-term intensity mapping applications, both the signal covariance and the foreground covariance should be assumed to be unknown. Thermal noise of the instrument can be measured well and is the only prior input to the estimator. We will develop expressions for generic \mathbf{Q}_α that will be tuned to be more robust to these unknowns.

With the choice $b_\alpha = \text{Tr}(\mathbf{N} \mathbf{Q}_\alpha)$, the expectation value of our estimator for the covariance model of Equation 3 is

$$\begin{aligned} \langle \hat{q}_\alpha \rangle &= \sum_\beta p_\beta \text{Tr}(\mathbf{C}_{,\beta} \mathbf{Q}_\alpha) \\ \langle \hat{\mathbf{p}} \rangle &= \mathbf{R} \langle \hat{\mathbf{q}} \rangle = \mathbf{R} \mathbf{M} \mathbf{p} = \mathbf{W} \mathbf{p}, \end{aligned} \quad (6)$$

where we have identified $\mathbf{M}|_{\alpha\beta} \equiv \text{Tr}(\mathbf{C}_{,\beta} \mathbf{Q}_\alpha)$ as a mixing matrix and $\mathbf{W} \equiv \mathbf{R} \mathbf{M}$ as the bandpower window functions. The origin of the mixing matrix is familiar from quadratic methods such as Hivon et al. (2002) and is due to the correlation of Fourier modes sampled over a finite area (or alternately not having an orthonormal basis in a restricted survey area).

The estimator $\hat{\mathbf{p}}$ is then a window \mathbf{W} that weights several modes of the true, underlying signal covariance amplitudes \mathbf{p} . The matrix \mathbf{M} is fully dictated by the estimator, but \mathbf{R} must be chosen. One option is to pick \mathbf{R} so that $\mathbf{W}_{\alpha\alpha} = 1$, ensuring that \hat{p}_α is a unit multiple of p_α . This does not mean that $\langle \hat{p}_\alpha \rangle = p_\alpha$ because \hat{p}_α will generally be a combination of several band powers. Another choice is to pick R_α to give a

weighted average of band powers as a window function

$$\sum_{\beta} \mathbf{W}|_{\alpha\beta} = 1. \quad (7)$$

This constraint does not fully specify \mathbf{R} , but we can impose an additional constraint for simplicity that \mathbf{R} is diagonal and $\mathbf{R}|_{\alpha\alpha} = R_{\alpha}$. This matrix normalizes each band power but does not apply any sense of decorrelation, which is developed as a final step in Section 6.3. In this case, the full estimator is

$$\hat{p}_{\alpha} = R_{\alpha}(\mathbf{x}^T \mathbf{Q}_{\alpha} \mathbf{x} - b_{\alpha}), \quad (8)$$

and the constraint in Equation 7 fixes

$$R_{\alpha} = \left[\sum_{\beta} \text{Tr}(\mathbf{C}_{,\beta} \mathbf{Q}_{\alpha}) \right]^{-1}. \quad (9)$$

We will refer to \mathbf{R} as the “rote” normalization because it handles any arbitrary multipliers of the quadratic data combination. For example, if we scale all the estimators \mathbf{Q}_{α} by γ , then $R_{\alpha} \propto 1/\gamma$, correcting for the normalization. Pencil-beam surveys and slabs (one beam wide in a slice many beams long) will have significant mode coupling and may warrant a different type of analysis. Here we consider application to map regions with essentially uniform sky coverage, such as ACT (Dünner et al. 2013) and SPT (Schaffer et al. 2011).

3.3. The Optimal Estimator and Its Normalization

So far, \mathbf{Q}_{α} has organized a generic quadratic combination of the data. In this section, we review the formally optimal estimator and the procedure it implies for analyzing the data. The optimal estimator provides a good starting point for constructing estimators, but subsequent statements will leave general \mathbf{Q}_{α} rather than assume optimality. The covariance of the \hat{q} values in Equation 4 is

$$\text{Cov}(\hat{q}_{\alpha}, \hat{q}_{\beta}) = \text{Tr}(\mathbf{C} \mathbf{Q}_{\alpha} \mathbf{C} \mathbf{Q}_{\beta}). \quad (10)$$

Minimizing the covariance results in

$$\mathbf{Q}_{\alpha} = \frac{\mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1}}{\text{Tr}(\mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1} \mathbf{C}_{,\alpha})}, \quad (11)$$

where the normalization is fixed by the Lagrange multiplier that forces the estimator to be unbiased.

Based on the previous section, the mixing matrix of the optimal estimator is

$$\mathbf{M}|_{\alpha\beta} = \text{Tr}(\mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1} \mathbf{C}_{,\beta}), \quad (12)$$

which is just the Fisher matrix \mathbf{F} of the estimator (e.g. Tegmark et al. (1997)).

This is the starting point for more robust estimators. Note that \mathbf{Q}_{α} is only optimal when \mathbf{C} is known exactly. When it is not known perfectly, the estimator is generally suboptimal but can be constructed to be unbiased.

3.4. Optimal Estimation Procedure

To get intuition for the term $\mathbf{C}_{,\alpha}$, note that in CMB analysis it is the outer product of spherical harmonics

$$\mathbf{C}_{,\alpha} = \mathbf{y}_{\alpha} \mathbf{y}_{\alpha}^T \text{ where } \mathbf{y}_{\alpha}|_i = Y_{\ell m(\alpha)}(\hat{\mathbf{r}}_i). \quad (13)$$

Hence, for the CMB, \hat{q}_{α} has the interpretation of the amplitude of the spherical harmonic transform of the weighted

maps, reinforcing \hat{q} as a pseudo-power as in Hivon et al. (2002).

The operation $\mathbf{C}_{,\alpha}$ is more complex in intensity mapping surveys because of the relation between the geometry of the observations and the desired 3D power spectrum. \mathbf{x} represents the calibrated intensity as a function of angular pointing and frequency and can be interpolated onto a cube in units of $h^{-1}\text{Mpc}$ using a fiducial cosmology. Here angular slices translate approximately into spatial slices, and frequencies translate into distances to redshifts. The relation between redshift and distance is nonlinear, and the low-frequency end of the survey represents a larger spatial area than the high-frequency side. The survey is therefore approximately a truncated pyramid in comoving Cartesian coordinates. Any areas outside the pyramid can be given zero weights in the larger coordinate volume. If the angular region is large enough, there will be non-flat sky curvature in the spatial slices. Through these factors, there is not a 1-to-1 translation between observation and comoving Cartesian coordinates, and interpolation inevitably leads to loss in fidelity. Consider linear interpolation schemes that can be represented as $\mathbf{P}\mathbf{x}$, where \mathbf{P} moves a map \mathbf{x} to Cartesian coordinates. It is not generally invertible, but so long as the discretization is fine enough, the observed $\theta - \nu$ space maps can be translated to Cartesian and back again with little loss.

Once in a grid of constant Cartesian dimensions, $\mathbf{C}_{,\alpha}$ can be understood as an operation that takes the 3D Fourier transform of both maps (Dillon et al. 2013) and bins the k -vectors into annuli in a range of $k_{\perp,\alpha}$ to $k_{\perp,\alpha+1}$ and $k_{\parallel,\alpha}$ to $k_{\parallel,\alpha+1}$ that define the range of the 2D band power p_{α} . Let \mathbf{K} be the linear Fourier operation so that $\tilde{\mathbf{x}} = \mathbf{K} \mathbf{P} \mathbf{C}^{-1} \mathbf{x}$ is the Fourier transform of the data in Cartesian coordinates. Then the binning operation is mathematically

$$\hat{q}_{\alpha} = \sum_{\mathbf{k}} I_{\mathbf{k} \in A_{\alpha}} \tilde{\mathbf{x}}(\mathbf{k}) \tilde{\mathbf{x}}(\mathbf{k})^* / \sum_{\mathbf{k}} I_{\mathbf{k} \in A_{\alpha}}, \quad (14)$$

where $I_{\mathbf{k} \in A_{\alpha}}$ is the indicator that is 1 in the k -bin annulus A_{α} and 0 elsewhere. Figure 1 shows the number of 3D Fourier modes contributing to 2D band powers with logarithmic spacing. This operation performs no weighting, assuming that $\mathbf{C}^{-1} \mathbf{x}$ has noise isotropic in the k_{\perp} annulus. The estimator has the form of an inner product $\tilde{\mathbf{x}}^T \mathbf{B}_{\alpha} \tilde{\mathbf{x}}$ where \mathbf{B}_{α} performs the binning of 3D Fourier cells. Combined, $\mathbf{C}_{,\alpha} = \mathbf{P}^T \mathbf{K}^T \mathbf{B}_{\alpha} \mathbf{K} \mathbf{P}$ and can be understood as taking the Fourier transform of both maps in Cartesian coordinates and then binning onto the band power α . This can be easily parallelized across band powers by using the same Cartesian space conversion and Fourier transform. Note that \mathbf{C}^{-1} is applied to the maps on both sides in observing coordinates of θ and ν . It will remain natural to discuss covariance in those dimensions rather than Cartesian coordinates because contamination naturally lives along ν .

The matrix calculations throughout this paper are skeletons that put the proper form to procedures implemented in software. Putting together what we have so far, $\mathbf{x}^T \mathbf{Q}_{\alpha} \mathbf{x}$ can be written as a numerically convenient procedure:

1. Weight both maps by their inverse covariance in observing coordinates (and remove the mean map if needed).
2. Translate observing to Cartesian coordinates.
3. Calculate the fast Fourier transform to both sides of the quadratic product.




Figure 1. Number of 3D Fourier modes contributing to a given band power in the GBT-wide survey. The survey region is $7^\circ \times 4.3^\circ$ from 700 to 900 MHz in 256 spectral bins. Toward increasing k_\perp and k_\parallel , the number of modes grows quadratically and linearly, respectively. The survey volume encompasses several fundamental Fourier modes that appear as isolated bands. In subsequent plots, we ignore these lone harmonics. Even though the number of available modes rapidly increases toward high k , information from $k_\perp > 0.4 \text{ hMpc}^{-1}$ is almost entirely suppressed by the GBT beam.

4. Bin onto a band power.

When considering one Fourier mode, the binning operation is separable (Liu et al. 2014) as $\mathbf{B}_\alpha = \mathbf{D}^T \mathbf{D}$. In this case, it is useful to derive the map in observing coordinates that has information about a Fourier mode. Doing this requires transforming through \mathbf{P} and back through \mathbf{P}^{-1} , which we assume to exist even though information could generally be lost in the repixelization. This “filter” for the modal information in a map becomes $\mathbf{x}'_\alpha = \mathbf{P}^{-1} \mathbf{K}^T \mathbf{D}_\alpha \mathbf{K} \mathbf{P} \mathbf{x}$. The map \mathbf{x}'_α is the Fourier component α of the noise-weighted input map, and the complete quadratic estimator is simply $\mathbf{x}_\alpha'^T \mathbf{x}'_\alpha$. In later sections, we will estimate rules of thumb for the impact of foreground operations on signal by considering these simple forms of inner products.

3.5. Estimating the Normalization

There is complete freedom in choosing \mathbf{R} in Equation 6. For arbitrary \mathbf{R} , $\mathbf{W} = \mathbf{R}\mathbf{M}$ translates theoretical expectations \mathbf{p} to the same space as the measured data $\hat{\mathbf{p}}$. If \mathbf{R} is invertible, there is no information loss, and the choice of \mathbf{R} relates to the presentation of data. For example, we argued that the choice of Equation 9 corrected for some multiplier γ in the estimator, but there is no reason it needs to. The theoretical $P(k)$ could be compared to the data through the same pipeline with the arbitrary γ retained, effectively comparing to $\gamma P(k)$. A non-invertible \mathbf{R} would be a poor choice because information is lost, but any information that does get through would permit a comparison of theory and observation.

The choice of \mathbf{R} is especially relevant to intensity mapping because signal is attenuated by the beam and foreground down-weighting. One outlook is that the theory $P(k)$ should have the same beam and foreground treatment applied as in the real data and then be compared to the “raw” experimental band powers, without making any effort to correct the observed band powers. This makes it difficult to intercompare experiments with different foreground properties and beams, or even the same experiment with different foreground treatments. We prefer to bring the measured $\hat{\mathbf{p}}$ into the same footing as the theoretical band powers \mathbf{p} from $P(k)$. This means

correcting for the beam (Section 3.6) and foreground cleaning (Section 4.2).

Normalizations take the form $\text{Tr}(\mathbf{C}_{\beta} \mathbf{Q}_\alpha)$ (see Equation 9). The previous section described a software procedure to calculate $\mathbf{x}^T \mathbf{Q} \mathbf{x}$, but not the trace with an arbitrary matrix. We would like to estimate the normalization using the same well-defined pipeline that calculates $\mathbf{x}^T \mathbf{Q} \mathbf{x}$, rather than develop a separate matrix operation. This reuse reduces the complexity of software, enforces consistency through common pipelines, and provides a convenient numerical implementation.

$\sum_\beta \text{Tr}(\mathbf{C}_{\beta} \mathbf{Q}_\alpha)$ is the expectation value of band powers when $p_\beta = 1$, so that we can estimate R_α in Monte Carlo by finding the mean $\hat{q}_\alpha = \mathbf{x}_{p=1}^T \mathbf{Q}_\alpha \mathbf{x}_{p=1}$ of white noise in the input $\mathbf{x}_{p=1}$ that is drawn from a covariance where $p_\beta = 1$. In other words, \mathbf{R} ensures that unit power in is unit power out. Monte Carlo estimation is efficient because there are typically many more map pixels than estimated band powers, so a single Monte Carlo involves significant averaging. The number of samples varies with estimator, but for GBT-wide, several hundred were typically sufficient.

Rather than lump the normalization into a single factor that is estimated with simulations, we prefer to partition the rote normalization, effects of beam convolution (Section 3.6), and foreground down-weighting as separate operations (Section 5.3). The rote normalization will ensure that the power spectral estimator of the spatially weighted maps properly recovers inputs. The foreground transfer function accounts for attenuation of the cosmological signal by foreground down-weighting. Section 4 develops foreground treatment that is naturally partitioned into operations that avoid contaminated modes and weigh the survey based on thermal noise. Finally, the beam transfer function accounts for differences between the measured band powers and the inputs due to convolution by the instrumental response. We develop the beam transfer function first because it is simplest.

3.6. Impact of the Instrumental Beam

Diffraction limits the resolution of single-dish (aperture) and interferometer instruments (baseline). We will only consider the case of the single-dish instrument with an axially symmetric beam. Beams that are not axially symmetric have a non-isotropic impact on the data in k -space. Also, typical surveys cover a region at several parallactic angles, so a beam that is not axially symmetric does not correspond to the stationary convolution across the map. When the beam convolution applies uniformly across the map, it is a multiplication in Fourier space, so let a particular Fourier mode be modulated as $\mathbf{C}_{\alpha}^B = B_\alpha \mathbf{C}_\alpha$. Throughout, the 2D band powers α combine the annulus in k_x and k_y at constant $k_\perp = \sqrt{k_x^2 + k_y^2}$, so B_α already implicitly assumes an axisymmetric beam. The expectation value of the estimator is

$$\langle \hat{q}_\alpha \rangle = \sum_\beta B_\beta p_\beta \text{Tr}(\mathbf{C}_{\beta} \mathbf{Q}_\alpha) = \mathbf{M} \mathbf{B} \mathbf{p} \quad (15)$$

$$\langle \hat{\mathbf{p}} \rangle = \mathbf{R}_B \langle \hat{\mathbf{q}} \rangle = \mathbf{R}_B \mathbf{M} \mathbf{B} \mathbf{p} = \mathbf{W}_B \mathbf{p}. \quad (16)$$

The window function $\mathbf{W}_B \equiv \mathbf{R}_B \mathbf{M} \mathbf{B}$ includes the effect of the beam, and \mathbf{R}_B reflects that we will want a different normalization due to the impact of the beam. The convolution acts on the underlying band powers, which are then mixed under \mathbf{M} through the estimator, identically to, e.g. Hivon et al.

(2002). In a context like the optimal quadratic estimator, where $\mathbf{Q}_\alpha = \mathbf{C}^{-1}\mathbf{C}_{,\alpha}^B\mathbf{C}^{-1}$, the estimator itself scales with the beam. In this case, \mathbf{M} also scales as B_α and the rote normalization Equation 9 scales as B_α^{-1} , so the effect cancels. We will neglect the impact of the beam on the estimator, assuming fixed \mathbf{Q}_α for simplicity.

Choose to keep the same rote normalization of the quadratic estimator as Equation 9 but extend the transformation \mathbf{R}_B to include a diagonal “beam transfer function” correction for the beam \mathbf{T}_B^{-1} , which is $\mathbf{R}_B \equiv \mathbf{T}_B^{-1}\mathbf{R}$. The expectation value is then

$$\langle \hat{\mathbf{p}} \rangle = \mathbf{T}_B^{-1}\mathbf{R}\mathbf{M}\mathbf{B}\mathbf{p} = \mathbf{W}_B\mathbf{p}. \quad (17)$$

The transfer beam function can be determined by enforcing $\sum_\beta \mathbf{W}_B|_{\alpha\beta} = 1$. Again choose a diagonal $T_\alpha^B = \mathbf{T}_B|_{\alpha\alpha}$ so that

$$(\mathbf{T}_\alpha^B)^{-1}R_\alpha\bar{B}_\alpha \sum_\beta \text{Tr}(\mathbf{C}_{,\beta}\mathbf{Q}_\alpha) = 1 \quad (18)$$

where

$$\bar{B}_\alpha = \sum_\beta \text{Tr}(\mathbf{C}_{,\beta}\mathbf{Q}_\alpha)B_\beta / \sum_\beta \text{Tr}(\mathbf{C}_{,\beta}\mathbf{Q}_\alpha). \quad (19)$$

This is the ratio of simulations with beam convolution to those without, which can be estimated using the same pipeline as the data. Plugging in the value from the rote normalization, Equation 9 yields $T_\alpha^B = \bar{B}_\alpha$. With the choices here, the transfer function is the mixing-weighted beam. Figure 2 shows the beam transfer function estimated from simulations. The full estimator has become

$$\hat{\mathbf{p}}_\alpha = \bar{B}_\alpha^{-1}R_\alpha(\mathbf{x}^T\mathbf{Q}_\alpha\mathbf{x} - b_\alpha). \quad (20)$$

This follows Tegmark (1997) rather than Hivon et al. (2002), where the data are multiplied by \mathbf{M}^{-1} and then the beam is treated. Our model throughout is to normalize the band powers (including the effect of beam and foreground) before decorrelating. Because of the effects of spatial and spectral masking and foreground deweighting, the mixing matrix of an intensity survey is not as easily calculable as the CMB. Here decorrelation is a final step related to display of the data and is based on simulations (Section 6.3) rather than a closed-form calculation.

The beam and its sidelobes generally broaden toward lower frequencies. As the sidelobes expand over the spatially varying continuum foreground structure, they can produce spectral structure. This is analogous to the wedge phenomenon in interferometers (see, e.g. Liu et al. (2014)). In a single-dish setting with approximately uniform map coverage, we can convolve all maps to a common resolution using the beam model. Figure 2 has no structure in k_\parallel because of this operation.

4. CLEANING KNOWN FOREGROUND MODES

The estimator developed so far has referred to general signal and noise covariances. Thermal noise clearly belongs in the noise covariance \mathbf{N} in Equation 3. Foregrounds are more challenging to attribute. Like thermal noise, they additively bias the power spectral estimates. Unlike thermal noise, they are not known accurately in advance, so the power spectral bias cannot be simply subtracted based on a model. They are also unavoidable astronomical signals, while thermal noise biases can be avoided by calculating the cross-power between subseasons with uncorrelated noise (as we will do in




Figure 2. Transfer function of the GBT beam. This is derived as the ratio of signal simulations with beam to those without. It is analogous to B_ℓ in CMB analysis, and here it clearly acts along the spatial (k_\perp) directions. Chromatic aspects of the beam are eliminated by convolving the maps to a common resolution, at the low-frequency end, with FWHM $\sim 0.3^\circ$.

Section 6.1). The perspective we take follows Switzer et al. (2013) and Dillon et al. (2014) by including any additive biases from residual foregrounds in the final power spectrum. This assumes that any foregrounds that can be modeled are down-weighted or subtracted, and residuals are, by definition, not possible to model and subtract as a noise bias. Residual foregrounds are indistinguishable from signal.

If foregrounds were drawn from a Gaussian distribution and known in advance, their covariance would be a complete description. This covariance would enter \mathbf{C}^{-1} and be down-weighted in the maps. In reality, we lack detailed knowledge of how a particular experiment will respond to bright foregrounds in a particular region of the sky. Further, both extragalactic fluctuations and the galaxy are non-Gaussian, so \mathbf{C}^{-1} will not properly down-weight the contaminated modes. A more conservative choice is to admit some generally contaminated modes that are fully projected out of the data rather than subjected to a more nuanced weight. These “avoided” modes are easy to include in the quadratic estimator.

Take the data to be true sky signal \mathbf{x}_{sig} plus some set of normalized modes \mathbf{Z} multiplied by amplitudes \mathbf{a} . The form of the quadratic estimator that we seek is then

$$\hat{q}_\alpha = (\mathbf{x}_{\text{sig}} + \mathbf{Z}\mathbf{a})^T\mathbf{Q}_\alpha(\mathbf{x}_{\text{sig}} + \mathbf{Z}\mathbf{a}) - b_\alpha, \quad (21)$$

where estimators are orthogonal to \mathbf{Z} , or $\mathbf{Q}_\alpha\mathbf{Z} = 0$. This choice forces \hat{q}_α to be uninfluenced by the modes in the matrix \mathbf{Z} . The quadratic estimator can be derived in the usual way of minimizing the covariance subject to the constraint that the result is unbiased (a normalization). In addition, a Lagrange multiplier can force $\mathbf{Q}_\alpha\mathbf{Z} = 0$. Appendix B4 of Tegmark et al. (1998) shows that the optimal

$$\mathbf{Q}_\alpha = \mathbf{\Pi}^T\mathbf{C}^{-1}\mathbf{C}_{,\alpha}\mathbf{C}^{-1}\mathbf{\Pi}, \quad (22)$$

$$\mathbf{\Pi} \equiv \mathbf{1} - \mathbf{Z}(\mathbf{Z}^T\mathbf{C}^{-1}\mathbf{Z})^{-1}\mathbf{Z}^T\mathbf{C}^{-1}. \quad (23)$$

Note that $\mathbf{\Pi}$ is idempotent $\mathbf{\Pi}\mathbf{\Pi} = \mathbf{\Pi}$ and so is a projection of the signal onto a space orthogonal to the modes \mathbf{Z} .

There are several equivalent viewpoints on foreground cleaning. Above, we have presented a constrained quadratic estimator that is orthogonal to contaminated modes. A Bayesian approach would jointly estimate contaminated modes with the signal modes in the map, but then marginal-

ize over contaminated modes to recover the signal. This marginalization is functionally equivalent to the orthogonal estimator (Seljak 1998). Both of these outlooks treat the contamination as additional “modes” to estimate, much like the signal. An alternative is to lump the contaminated modes with the noise covariance that is handled by the \mathbf{C}^{-1} operation. This too is connected to the above approach through the Woodbury identity

$$\mathbf{C}^{-1}\mathbf{\Pi} = \lim_{\sigma^2 \rightarrow \infty} (\mathbf{C} + \sigma^2 \mathbf{Z}\mathbf{Z}^T)^{-1} = (\mathbf{C} + \mathbf{F})^{-1}. \quad (24)$$

The estimator avoids the foreground modes \mathbf{Z} , so \mathbf{C} represents thermal noise, cosmological signal, and any residual foregrounds that are not explained by the modes \mathbf{Z} .

Again there is freedom in choosing \mathbf{C} to yield slightly less optimal but more robust estimators. We will continue to ignore the signal covariance contribution in the inverse noise weights. Until a high-significance signal detection is reached, signal covariance will be both subdominant and imprecisely known. We will assume that \mathbf{C} is purely thermal noise and is diagonal. \mathbf{C}^{-1} translates simply into weighting pixels by the inverse of the variance from thermal noise, which is related to the integration time and effective $T_{\text{sys}}(\nu)$ of the survey. This prescription allows \mathbf{C} to be precisely determined using survey properties and checked using the difference maps between subseasons (Section 6.1).

We can correct for the impact of foreground cleaning in the same scheme as the beam. Let the band power estimate be

$$\hat{p}_\alpha = (T_\alpha^B T_\alpha^F)^{-1} R_\alpha (\mathbf{x}^T \mathbf{Q}_\alpha \mathbf{x} - b_\alpha), \quad (25)$$

where T_α^B is the beam transfer function correction and R_α is the rote normalization, as before. We can form a foreground cleaning transfer function T_α^F using simulations that compare the ratio of the estimated band power α before and after foreground cleaning is applied to the data. More formally,

$$T_\alpha^F = \frac{\sum_\beta \text{Tr}(\mathbf{C}_{\beta,\beta} \mathbf{\Pi}^T \mathbf{C}^{-1} \mathbf{C}_{\alpha,\alpha} \mathbf{C}^{-1} \mathbf{\Pi})}{\sum_\beta \text{Tr}(\mathbf{C}_{\beta,\beta} \mathbf{C}^{-1} \mathbf{C}_{\alpha,\alpha} \mathbf{C}^{-1})}. \quad (26)$$

The transfer function must be measured in 2D k -space rather than 1D k shells because of the anisotropic nature of the cleaning. Foreground cleaning is not a real-space convolution like the beam, so is not fully described by a band power multiplier. The factors $(T_\alpha^B T_\alpha^F)^{-1} R_\alpha$ give the proper normalization but do not attempt to de-correlate the band powers. Section 6.3 describes how covariance simulations can be used to produce a final summary with independent errors per band power.

4.1. Line-of-sight Cleaning: Known Modes

The modes in \mathbf{Z} could contain both spatial and spectral information. Most of the foreground variance that distinguishes it from the signal is in the ν, ν' directions (Liu & Tegmark 2011), so \mathbf{Z} can be separated into spectral modes per line of sight and yield a separable projection $\mathbf{\Pi} = \mathbf{\Pi}_\nu \otimes \mathbf{1}$. Rather than identifying a handful of contaminated modes, it is convenient to work with a complete spectral foreground basis where only a handful of modes are projected out. Starting with a limited set of foreground spectral modes \mathbf{u}_i^f , a complete basis can be constructed using a Gram-Schmidt process. (Superscripts without a numerical value are used as labels rather than powers.) Assemble these complete basis vectors into columns of a matrix \mathbf{U}_f that projects onto a “spectral foreground” basis.

The matrix \mathbf{X} representation of the map (Section 3.1) has dimensions $N_\nu \times N_\theta$ and can be put in this basis as $\mathbf{U}_f^T \mathbf{X}$. This combination is a stack of maps of the amplitudes of spectral modes. The line-of-sight projection operation becomes

$$\mathbf{\Pi} \mathbf{x} = \text{vec}([\mathbf{1} - \mathbf{U}_f \mathbf{S} (\mathbf{U}_f^T \mathbf{C}_{\nu\nu'}^{-1} \mathbf{U}_f)^{-1} \mathbf{U}_f^T \mathbf{C}_{\nu\nu'}^{-1}] \mathbf{X}), \quad (27)$$

where we have added a diagonal selection matrix \mathbf{S} that is 1 for those modes that are subtracted and 0 for spectral modes that pass through. The inverse covariance in this case is $\mathbf{C}_{\nu\nu'}^{-1}$, which under the assumptions here is just the ν, ν' covariance of the thermal noise. This is a diagonal matrix derived from $T_{\text{sys}}(\nu)$. For notational simplicity, we will assume that the thermal noise is constant across all frequencies so that $\mathbf{C}_{\nu\nu'}^{-1} = \mathbf{1}$ and

$$\mathbf{\Pi} \mathbf{x} = \text{vec}((\mathbf{1} - \mathbf{U}_f \mathbf{S} \mathbf{U}_f^T) \mathbf{X}), \quad (28)$$

using orthonormality of the modes \mathbf{U}_f . This has the form of a simple projection of spectral modes. Experiments whose noise varies with frequency (or has masked frequencies) should use the $\mathbf{C}_{\nu\nu'}^{-1}$ -weighted form. Rather than carry around the full $\mathbf{\Pi}$, define the line-of-sight projection that acts on the left side of the map \mathbf{X} as

$$\mathbf{\Pi}_\nu \equiv \mathbf{1} - \mathbf{U}_f \mathbf{S} \mathbf{U}_f^T. \quad (29)$$

4.2. Cleaning Effectiveness and Direct Loss

In this section, we use the variance as a simple measure of the impact of foreground cleaning, rather than the more complex bandpower. The variance is directly related to the eigenvalue spectrum and modal structure of signal foregrounds and provides rules of thumb and intuition about the process of foreground cleaning. Let the ν, ν' covariance of a pure signal map \mathbf{X}_s be $\mathbf{C}_s = N_\theta^{-1} \langle \mathbf{X}_s \mathbf{X}_s^T \rangle$. Diagonalize the signal covariance as $\mathbf{C}_s = \mathbf{U}_s \mathbf{\Lambda}_s \mathbf{U}_s^T$, where $\mathbf{\Lambda}_s|_{ii} = \lambda_i^s$ and \mathbf{U}_s is made of the signal spectral eigenvectors \mathbf{u}_i^s . The signal variance is

$$\xi_s = N_\theta^{-1} \langle \mathbf{x}_s^T \mathbf{x}_s \rangle = N_\theta^{-1} \langle \text{Tr}(\mathbf{X}_s \mathbf{X}_s^T) \rangle \quad (30)$$

$$= \text{Tr}(\mathbf{U}_s \mathbf{\Lambda}_s \mathbf{U}_s^T) = \sum_k \lambda_k^s, \quad (31)$$

where the expectation value is over signal realizations, and we have identified the Frobenius trace $\text{Tr}(\mathbf{X}_s \mathbf{X}_s^T)$. This is the simplest quadratic estimator, corresponding to $\mathbf{Q} = \mathbf{1}$. Recall that the operation $\mathbf{x}'_\alpha = \mathbf{P}^{-1} \mathbf{K}^T \mathbf{D}_\alpha \mathbf{K} \mathbf{P} \mathbf{x}$ developed in Section 3.4 can filter all of the modal content of some band power α onto a new map \mathbf{x}'_α so that $\mathbf{x}_\alpha'^T \mathbf{x}'_\alpha$ is the full quadratic estimator.

Similar to the signal, pure foreground spectral covariance can also be decomposed into eigenmodes \mathbf{U}_f as $\mathbf{C}_f = N_\theta^{-1} \langle \mathbf{X}_f \mathbf{X}_f^T \rangle = \mathbf{U}_f \mathbf{\Lambda}_f \mathbf{U}_f^T$. After applying the cleaning $\mathbf{\Pi}_\nu$, the trace of the foregrounds becomes

$$\xi_{clean}^{fg} = N_\theta^{-1} \text{Tr}(\mathbf{\Pi}_\nu \mathbf{X}_f \mathbf{X}_f^T \mathbf{\Pi}_\nu^T) = \sum_{i \notin \text{cuts}} \lambda_i^f. \quad (32)$$

The variance of residual foregrounds in the map is the sum of the eigenvalues of the modes that were not projected out.

The foreground cleaning operation does not null discrete signal modes because signal and foreground have a different basis. A good measure of this signal loss is to apply foreground cleaning to a signal map ($\mathbf{\Pi}_\nu \mathbf{X}_s$) and find the cross-

variance with the input signal map \mathbf{X}_s ,

$$\xi_{\text{clean}} = N_{\theta}^{-1} \langle \text{Tr}(\mathbf{\Pi}_{\nu} \mathbf{X}_s \mathbf{X}_s^T) \rangle = \xi_s + \xi_{\text{direct}} \quad (33)$$

$$\xi_{\text{direct}} \equiv -N_{\theta}^{-1} \langle \text{Tr}(\mathbf{U}_f \mathbf{S} \mathbf{U}_f^T \mathbf{X}_s \mathbf{X}_s^T) \rangle, \quad (34)$$

where we have identified the signal covariance from Equation 31 and defined the loss of signal as a direct result of cleaning as ξ_{direct} , which will be negative. Hence, the signal variance in the cleaned map is the input signal variance minus some loss. We will show that this loss is the overlap of signal with the foreground modes that were projected out. Writing out the direct loss,

$$\xi_{\text{direct}} = - \sum_{j \in \text{cuts}, k} \lambda_k^s \langle [(\mathbf{u}_k^s)^T \mathbf{u}_j^f]^2 \rangle, \quad (35)$$

which is the overlap of the signal modes with the foreground modes that are subtracted. We can develop a simple expression by assuming that any foreground mode is uncorrelated with any random signal mode. In this case, spurious correlations scale as $(\mathbf{u}_k^s)^T \mathbf{u}_j^f \propto 1/\sqrt{N_{\text{res}, \nu}}$, where $N_{\text{res}, \nu}$ is the number of independent modes of the signal over the frequency range. For example, at the highest k in the box, this approaches N_{ν} , while for lower k there are relatively fewer spectral modes in the signal. If N_m modes are removed in the sum over $j \in \text{cuts}$,

$$\xi_{\text{clean}} = \xi_s + \xi_{\text{direct}} \sim \left(1 - \frac{N_m}{N_{\text{res}, \nu}}\right) \xi_s. \quad (36)$$

Detection of signal in the cleaned maps benefits as the number of signal degrees of freedom $N_{\text{res}, \nu}$ exceeds the number of foreground modes removed, N_m .

There will generally be some spurious correlation between foreground spectral modes and signal spectral modes. Under an assumption that the foregrounds are physically unrelated to the signal, there is no correlation to the signal, on average. Hence, spurious correlation is an issue of noise and signal loss rather than bias. (This can be violated when the sources of line radiation produce continuum emission, but that case is left for future work.)

In practice, we expect that intensity mapping experiments will need to measure at least some contaminated modes from the data themselves. In this case, the signal influences the foreground modes. We next show that this produces a net anti-correlation of the cleaned foregrounds with the signal, resulting in considerable bias on the largest scales in the map. This effect is familiar from ILC bias in CMB foreground cleaning (see, e.g. Efstathiou et al. (2009)), and in general of blind methods.

5. EMPIRICAL CLEANING

The instrument response to bright foregrounds will not follow a Gaussian distribution or be fully known in advance. Some examples of instrument response are passband calibration, chromatic beam response, polarization leakage, and calibration stability. Experiments will make a best effort at calibrating these effects, but any differences will modulate how the experiment observes bright foregrounds. Further, we would like to avoid assuming that the intrinsic spectrum of contaminants is known in advance.

A limited number of contaminated spectral modes can be estimated from the data themselves by finding the empirical ν, ν' covariance $\hat{\mathbf{C}}$ and its eigenvalue decomposition $\hat{\mathbf{C}} =$

$N_{\theta}^{-1} \mathbf{X} \mathbf{X}^T = \hat{\mathbf{U}} \hat{\mathbf{\Lambda}} \hat{\mathbf{U}}^T$. Unlike the previous section, we do not take the average over signal realizations through $\langle \rangle$, assuming instead that we only have access to one map \mathbf{X} that is the sum of foregrounds, signal, and noise. Writing $\hat{\mathbf{C}}$ in terms of these constituents,

$$\hat{\mathbf{C}} = N_{\theta}^{-1} (\mathbf{X}_f + \mathbf{X}_s + \mathbf{X}_n) (\mathbf{X}_f + \mathbf{X}_s + \mathbf{X}_n)^T. \quad (37)$$

Initially neglect the noise contribution so that

$$\begin{aligned} \hat{\mathbf{C}}_{s+f} &= N_{\theta}^{-1} (\mathbf{X}_f \mathbf{X}_f^T + \mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T + \mathbf{X}_s \mathbf{X}_s^T) \\ &= \hat{\mathbf{C}}_f + \mathbf{C}_{\Delta}, \end{aligned} \quad (38)$$

where $\mathbf{C}_{\Delta} \equiv N_{\theta}^{-1} (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T + \mathbf{X}_s \mathbf{X}_s^T)$, and $s + f$ denotes the fact that the covariance contains both signal and foregrounds. Thus, the eigenvalues that we estimate from $\hat{\mathbf{C}}$ likely are dominated by foreground $\hat{\mathbf{C}}_f$, but also perturbed by the signal itself (and noise) through \mathbf{C}_{Δ} . Write the updated eigenvalues of $\hat{\mathbf{C}}$ as

$$\mathbf{U}_{s+f} = \mathbf{U}_f + \mathbf{\Delta}. \quad (39)$$

The data-cleaning operation is then with respect to these perturbed eigenvectors, defining

$$\mathbf{\Pi}_{s+f}^{\nu} \equiv 1 - (\mathbf{U}_f + \mathbf{\Delta}) \mathbf{S} (\mathbf{U}_f + \mathbf{\Delta})^T. \quad (40)$$

This again projects a subset of contaminated modes from each line of sight. In comparison, the ideal cleaning would only remove foreground modes as

$$\mathbf{\Pi}_f^{\nu} \equiv 1 - \mathbf{U}_f \mathbf{S} \mathbf{U}_f^T. \quad (41)$$

There is not a general analytic theory describing the eigenvectors of the sums of matrices. It is possible to write a rank- N update of the eigenvectors that gives some intuition that the signal “rotates” foreground modes, but that direction does not yield analytic expressions. The next section develops expressions that are perturbative in small signal, and some analytic insights are possible. Ultimately, the properties of this cleaning need to be estimated numerically in transfer function simulations (Section 5.3).

5.1. Example: GBT Foreground Decomposition and Instrument Response

Figure 3 shows the eigenvalue spectrum of foregrounds and noise in the GBT-wide survey. After removing 10 modes, the foreground fluctuations in map space are suppressed by $\sim 10^3$; however, the covariance remaining in the map is spread in a tail of modes at lower amplitude. Figure 3 also shows the eigenvalue spectrum of the ν, ν' covariance of the difference between maps of subseasons (Section 6.1), which cancels constant astronomical signal. Before taking the map difference, we recalibrate each line of sight to isolate variation in spectral shape rather than amplitude. Beyond the brightest ~ 10 modes, instrumental noise and time-varying spectral response become increasingly significant. Note that the eigenvalue spectrum of the sum of covariances (noise, signal, and foreground) is generally not the sum of the spectra of their respective covariances. Hence, the the total eigenvalue spectra cannot be rigorously decomposed into the sum of signal, foreground, and noise parts. A well-designed experiment should have the majority of foreground covariance explained by a few modes. Section 5.4 describes general conditions in which signal can be recovered well.




Figure 3. Square root of the eigenvalue spectrum of the ν, ν' covariance of the GBT-wide field, normalized to one for the largest foreground mode. The square root converts variance to rms temperature fluctuations in the map. The solid line shows the spectrum of the ν, ν' covariance of the input maps, while the dashed line shows the spectrum of the ν, ν' covariance of the difference between maps of subseasons. The difference of subseason maps removes astronomical signal and foregrounds that are common across observations and isolates any time-varying noise or instrumental systematics (Section 6.1). The largest 10^3 of the rms can be explained by ~ 10 modes. Beyond that peak is a plateau of much smaller modes, which are increasingly dominated by noise. See Figure 4 for the first five eigenvectors. The downturn at high mode number reflects the impact of the finite number of spectrometer channels, with some fraction cut due to RFI contamination.

Figure 4 shows the largest five eigenvectors of the GBT-wide field. The largest mode is effectively the mean synchrotron emission across the survey. Commonly proposed smooth spectral functions (polynomials, power laws, foreground model eigenvectors) fail to explain both the small glitches and the overall, non-power-law structure. These residuals would still be hundreds of times larger than the signal. It is possible to recalibrate the data such that the largest mode is spectrally smooth, based on the largest eigenvector. However, this does not improve the overall prospects for foreground deweighting. By the construction of the eigenvectors, the higher modes are specifically orthogonal to the variations of the first. Even if the mean synchrotron were forced to be smooth, it would have no impact on the remaining, independent contaminant modes. The brightest modes are very poorly described by standard series of orthogonal polynomials. At high enough order, a complete polynomial basis could explain any of these foregrounds; however, it would also explain and project out the signal. By discovering contaminated modes in the data themselves, we remove variance surgically-leaving the most remaining degrees of freedom for signal to transfer through.

Most of the structure in Figure 4 is due to instrumental response. This section reviews the beam and spectral calibration process employed with the GBT data, emphasizing aspects relevant to contaminant modes. Complete details can be found in Masui (2013).

A broadband noise source injects power at the feed point and acts as a stable flux reference. We switch this noise source with a rapid 64 ms period that allows the measured calibration signal to be uncontaminated by sky signal and RFI. The data from a single scan (which for the wide-field data used here is 2 minutes in length) are referenced to the mean amplitude of the noise calibrator signal. The noise calibrator itself is assumed to be perfectly stable. The noise calibration (a time trans-

fer standard) is then referenced using a collection of scans of well-characterized, bright, unpolarized point sources such as 3C 48, 3C 295, and 3C 147. This procedure is performed independently in each spectral channel and for the power from the X and Y polarized antennae. Separately calibrating X and Y signals mitigates leakage of polarization into the summed unpolarized intensity. Analysis of the point source data suggests that the primary source of polarization leakage on boresight is due to a difference in gain between the X and Y antennae.

The above procedure achieves 0.5% uncertainty per band over 1 minute of integration across 84 total hours of integration in the GBT wide field survey. The data used here assume that the noise calibrator is completely stable over the observation. Similar systems (Bryerton 2011) are known to vary at the $\sim 1\%$ level. Future work must characterize the stability of noise calibrator and the covariance of its spectrum. Variations in the calibrator’s spectral structure could result in a proliferation of contaminant degrees of freedom, while common mode variations can be more benign. The derived calibration in our GBT data varied at the 1%-level and may be partly attributable to the calibration source rather than receiver stability.

The GBT prime-focus beams have significant polarization leakage off boresight because of the off-axis design. Polarization leakage between linear polarization and total intensity is of order 5% of the primary gain. The polarization leakage can cause spectral structure in the foreground (and thus additional degrees of freedom) in two ways: (1) Faraday rotation of polarized synchrotron emission can vary across lines of sight and mix to frequency structure in the intensity spectrum and (2) leakage from polarization to intensity has its own spectral structure due to the instrument. The latter is conclusively observed in the third mode of Figure 4.

Mitigation of polarization to intensity leakage is a subject for future work. A promising avenue is that the Mueller leakage beams are observed to be approximately odd functions about the boresight axis. Smoothing the maps above the beam scale tends to suppress this leakage. Spatial smoothing decreases the number of degrees of freedom that the cosmological signal can exercise. We will argue below that there is a significant penalty for fewer signal degrees of freedom because of spurious correlation of the signal and foregrounds. The most promising approach may be to pursue the “deprojection” approach used in the BICEP CMB polarization analysis (Takahashi et al. 2008; Kaufman et al. 2014) where contamination from leakage could be removed in the time domain data using a model of the response and polarized emission.

5.2. Perturbative Expansion for Signal-foreground Correlations

Write out the cleaned map for small signals as

$$\begin{aligned} \mathbf{X}_{\text{clean}} &= \mathbf{\Pi}_{\text{s+f}}^{\nu} (\mathbf{X}_{\text{s}} + \mathbf{X}_{\text{f}}) \\ &= [1 - (\mathbf{U}_{\text{f}} + \mathbf{\Delta})\mathbf{S}(\mathbf{U}_{\text{f}} + \mathbf{\Delta})^T](\mathbf{X}_{\text{s}} + \mathbf{X}_{\text{f}}) \\ &\approx \mathbf{\Pi}_{\text{f}}^{\nu} \mathbf{X}_{\text{s}} + \mathbf{\Pi}_{\text{f}}^{\nu} \mathbf{X}_{\text{f}} - \mathbf{\Delta}\mathbf{S}\mathbf{U}_{\text{f}}\mathbf{X}_{\text{f}} - \mathbf{U}_{\text{f}}\mathbf{S}\mathbf{\Delta}^T \mathbf{X}_{\text{f}}. \end{aligned} \quad (42)$$

Here $\mathbf{\Pi}_{\text{f}}^{\nu} \mathbf{X}_{\text{s}}$ and $\mathbf{\Pi}_{\text{f}}^{\nu} \mathbf{X}_{\text{f}}$ are signal and foreground cleaned by pure foreground modes and are analogous to the direct loss in Equations 35 and 32, respectively. Combine the remaining terms as

$$\tilde{\mathbf{X}}_{\text{s}} \equiv -\mathbf{\Delta}\mathbf{S}\mathbf{U}_{\text{f}}\mathbf{X}_{\text{f}} - \mathbf{U}_{\text{f}}\mathbf{S}\mathbf{\Delta}^T \mathbf{X}_{\text{f}}. \quad (43)$$

The combination $\mathbf{\Pi}_{\text{f}}^{\nu} \mathbf{X}_{\text{f}} + \tilde{\mathbf{X}}_{\text{s}}$ represents residual foregrounds in the map after the empirical foreground cleaning is applied.




Figure 4. First five eigenvectors of the ν, ν' covariance in the GBT-wide field, descending in amplitude from the top. The largest mode is the mean synchrotron emission across the map. Despite best efforts at calibrating, there is visible small structure and non-power-law behavior. The second-largest mode is due to the chromatic beam response and can be effectively nulled by convolving the maps to a common resolution. The third mode has amplitudes that agree well spatially with polarized emission in the survey area, and the frequency structure corresponds approximately to the frequency dependence of polarization to intensity leakage. Beyond this mode, no clear instrumental response systematics can be discerned. Possibilities are calibration instability or variations in rotation measure appearing through polarization leakage. To reach the level of the signal, ~ 30 modes need to be removed. The spectral modes are very poorly approximated by a series of smooth polynomials.

In $\Pi_f^\nu \mathbf{X}_f$, both Π_f^ν and \mathbf{X}_f are assumed to be unrelated to the signal, while the second term $\tilde{\mathbf{X}}_s$ is a scaled version of the perturbations to the foreground modes due to the signal, Δ . Considering the minus signs, these residuals are anticorrelated with the signal. We will show that this term biases the band powers in a way that does not average to zero over many signal realizations, analogous to the ILC bias (e.g. Efstathiou et al. (2009)).

When estimating the impact of empirical foreground cleaning on the signal, it is insufficient to consider only the direct loss of the signal as $\Pi_{s+f}^\nu \mathbf{X}_s$. One must also include the impact of spurious correlations of the signal and foregrounds. These terms are only manifest if we simulate $\Pi_{s+f}^\nu (\mathbf{X}_s + \mathbf{X}_f)$, including the foregrounds.

As before, to get a rule of thumb, take the cross-correlation between the cleaned map and pure signal, as

$$\begin{aligned} \xi_{\text{clean}} &= N_\theta^{-1} \langle \text{Tr}(\mathbf{X}_{\text{clean}} \mathbf{X}_s^T) \rangle \\ &= N_\theta^{-1} \langle \text{Tr}(\Pi_f^\nu \mathbf{X}_s \mathbf{X}_s^T) \rangle + N_\theta^{-1} \langle \text{Tr}(\tilde{\mathbf{X}}_s \mathbf{X}_s^T) \rangle, \end{aligned} \quad (44)$$

where $\langle \rangle$ is over signal realizations. We have neglected the piece $\langle \text{Tr}(\Pi_f^\nu \mathbf{X}_f \mathbf{X}_s^T) \rangle$ because the foregrounds cleaned with pure foreground modes $\Pi_f^\nu \mathbf{X}_f$ have no expected correlation to the signal. Equation 34 describes the first term as being the signal and some direct loss. However, now that foreground modes are estimated from the map itself, the final term $N_\theta^{-1} \langle \text{Tr}(\tilde{\mathbf{X}}_s \mathbf{X}_s^T) \rangle$ describes the spurious average correlation of foreground residuals and the signal. Writing these out,

$$\xi_{\text{clean}} = \xi_s + \xi_{\text{direct}} + \xi_{\text{spur}}. \quad (45)$$

We now calculate ξ_{spur} in a perturbative limit and show that it significantly impacts the signal. At first order, the perturbed vectors \mathbf{U}_{s+f} are a linear combination of the pure foreground eigenvectors \mathbf{U}_f times a weight related to the perturbation. For small signals, the perturbation to the covariance is $\mathbf{C}_\Delta = N_\theta^{-1} (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T)$, related to the spuri-

ous correlation of signal and foreground over a finite patch of the sky with N_θ samples. According to first-order perturbation theory, the vectors δ_i in the perturbation matrix Δ of $\mathbf{U}_{s+f} = \mathbf{U}_f + \Delta$ are

$$\delta_i = \sum_j \frac{(\mathbf{u}_i^f)^T \mathbf{C}_\Delta \mathbf{u}_j^f}{\lambda_i^f - \lambda_j^f} \mathbf{u}_j^f. \quad (46)$$

Recall that the foreground eigenvectors $\{\mathbf{u}_i^f\}$ are assumed to have full rank and are the columns of \mathbf{U}_f . Through some algebra (Appendix A),

$$\xi_{\text{spur}} = -N_\theta^{-1} \left\langle \sum_{\substack{i \in \text{cuts} \\ j \notin \text{cuts}}} \frac{[(\mathbf{u}_i^f)^T (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T) \mathbf{u}_j^f]^2}{\lambda_i^f - \lambda_j^f} \right\rangle \quad (47)$$

The terms $\mathbf{X}_f \mathbf{X}_s^T$ describe spurious correlation of signal and foreground and fluctuate about zero across signal realizations in the $\langle \rangle$ average. The fact that this expression is squared inside the $\langle \rangle$ means that there is a net anticorrelation of the foreground residuals with the signal itself. This anticorrelation depends in detail on the particular foreground \mathbf{X}_f in the map region. Section 5.3 considers this numerically for the GBT-wide survey, but it is convenient to have an analytic rule of thumb for the signal that is attenuated as a function of number of modes removed N_m . Under assumptions similar to Equation 36, Appendix A finds

$$\xi_{\text{clean}} = \left(1 - \frac{N_m}{N_{\text{res}, \nu}}\right) \left(1 - \frac{N_m}{N_{\text{res}, \theta}}\right) \xi_s, \quad (48)$$

where $N_{\text{res}, \nu}$ and $N_{\text{res}, \theta}$ are the number of modes in the frequency (k_\parallel) and angular (k_\perp) directions, respectively. The general scaling $1 - N_m/N_{\text{res}, \nu}$ is intuitive: if there are only 10 spectral degrees of freedom in the signal and we need to remove 10 spectral modes, no signal remains. The dependence on the angular component is less intuitive because our cleaning operation acts entirely in the frequency direction. Indeed, if the foreground modes were taken as given, the signal loss scales as Equation 36. The dependence on the number of angular modes $N_{\text{res}, \theta}$ arises as a by-product of measuring the foreground spectral modes from the map itself. The covariance $\hat{\mathbf{C}} = N_\theta \mathbf{X} \mathbf{X}^T$ measures foreground spectral modes against a limited number of signal realizations. In general, if the same foreground mode is observed against many different signal realizations, the spurious correlation averages down.

Equation 48 is only a rule of thumb, and simulations are needed to effectively measure the number of resolution elements $N_{\text{res}, \nu}$ and $N_{\text{res}, \theta}$. These quantities are not simply N_ν and N_θ , the number of spectral and angular pixels in the survey. Instead, $N_{\text{res}, \nu}$ and $N_{\text{res}, \theta}$ roughly relate to the number of modes in the signal at the \mathbf{k} scales of interest. Signal components at the lowest k_\perp of the survey have large wavelengths and few $N_{\text{res}, \theta}$ modes in the survey volume. The spurious correlation of signal and foreground does not average down over many signal modes, and so most of the large-scale signal is lost. The impact of foreground cleaning on the signal depends strongly on both k_\perp and k_\parallel , and so must be treated as a 2D transfer function.

From the point of view of survey design, one wants to maximize the number of resolution elements $N_{\text{res}, \nu}$ and $N_{\text{res}, \theta}$. The beam size fixes the ultimate sensitivity at high k_\perp , and having high $N_{\text{res}, \theta}$ translates into large surveys with many

beam spots. In general, intensity mapping surveys will need to cover larger areas than Fisher estimates from thermal noise suggest, owing to the fact that potentially many degrees of freedom in the data need to be used to estimate foregrounds.

5.3. Estimating the Signal Loss Transfer Function

The previous section argued that when contaminated modes are discovered from the data themselves, residual foregrounds will be anticorrelated with the signal (on average). The means that to assess signal attenuation due to foreground cleaning, it is insufficient to apply the cleaning to signal-only simulations and measure loss as in Section 3.6. Instead, the foreground cleaning should be applied to simulations of signal plus foregrounds. A reasonable approach is to add signal simulations \mathbf{x}_{sim} to the measured map itself, \mathbf{x} (which is assumed to be dominated by continuum foregrounds). One can then subtract the foreground power back from the band power, as

$$\hat{q}_\alpha|_{\text{out}} = [\mathbf{C}^{-1}\mathbf{\Pi}_{\text{s+f}}^\nu(\mathbf{x} + \mathbf{x}_{\text{sim}})]^T \mathbf{Q}_\alpha \mathbf{C}^{-1}\mathbf{\Pi}_{\text{s+f}}^\nu(\mathbf{x} + \mathbf{x}_{\text{sim}}) - (\mathbf{C}^{-1}\mathbf{\Pi}_f^\nu \mathbf{x})^T \mathbf{Q}_\alpha \mathbf{C}^{-1}\mathbf{\Pi}_f^\nu \mathbf{x}, \quad (49)$$

where the modes of the cleaning operation $\mathbf{\Pi}_{\text{s+f}}^\nu$ are determined from the map of real data plus simulations ($\mathbf{x} + \mathbf{x}_{\text{sim}}$) and $\mathbf{\Pi}_f^\nu$ is determined from only the real data, taken to define the foreground modes. Equation 49 represents the band power estimate of the remaining signal simulation, and the transfer function can be estimated as the ratio of this quantity to the input simulation band power weighted by the thermal noise of the survey, or

$$\hat{q}_\alpha|_{\text{in}} = (\mathbf{C}^{-1}\mathbf{x}_{\text{sim}})^T \mathbf{Q}_\alpha \mathbf{C}^{-1}\mathbf{x}_{\text{sim}}, \quad T_\alpha^F = \left\langle \frac{\hat{q}_\alpha|_{\text{out}}}{\hat{q}_\alpha|_{\text{in}}} \right\rangle. \quad (50)$$

This is analogous to Equation 26 but includes spurious correlation of signal and foreground. Cosmological signals with poorly understood amplitudes (such as when a cross-power with a galaxy survey is not available) may need a range of signal simulation amplitudes to understand any sensitivity of the transfer function.

The optimal number of modes to remove is described next in Section 5.4. To be able to compare outcomes for different numbers of modes removed, the transfer function needs to be calculated for a range of scenarios. Especially for too few modes removed, the residual foreground variance may be significant. In this case, the estimate of $\hat{q}_\alpha|_{\text{out}}$ is noisy and Equation 50 requires an average over many realizations to converge satisfactorily.

Write out the cleaned map of Equation 43 in \mathbf{x} rather than \mathbf{X} notation as

$$\mathbf{x}_{\text{clean}} = \mathbf{\Pi}_{\text{s+f}}(\mathbf{x} + \mathbf{x}_{\text{sim}}) \approx \mathbf{\Pi}_f \mathbf{x} + \mathbf{\Pi}_f \mathbf{x}_{\text{sim}} + \tilde{\mathbf{x}}_{\text{sim}}, \quad (51)$$

where again $\tilde{\mathbf{x}}_{\text{sim}}$ is the critical piece of residual foregrounds that are anticorrelated with simulated signal. On average, the term $\mathbf{\Pi}_f \mathbf{x}$ is just producing variance and contains nothing related to the signal, so it can be subtracted. We can then form $\mathbf{\Pi}_{\text{s+f}}(\mathbf{x} + \mathbf{x}_{\text{sim}}) - \mathbf{\Pi}_f \mathbf{x} = \mathbf{\Pi}_f \mathbf{x}_{\text{sim}} + \tilde{\mathbf{x}}_{\text{sim}}$ to get just those pieces relevant to understanding how the signal acts under foreground subtraction. Signal loss can then be assessed in the cross-power between this cleaned signal map and the input signal

$$\hat{q}_\alpha|_{\text{loss}} = [\mathbf{\Pi}_{\text{s+f}}(\mathbf{x} + \mathbf{x}_{\text{sim}}) - \mathbf{\Pi}_f \mathbf{x}]^T \mathbf{C}^{-1} \mathbf{Q}_\alpha \mathbf{x}_{\text{sim}}. \quad (52)$$

We refer to this as a one-sided estimate of signal loss, because $\mathbf{\Pi}$ is only on one side of the quadratic estimator. The numer-

ator of the transfer function needs the expectation value of signal with foreground cleaning applied to both sides. We can rewrite this using a relation between trace and covariance as

$$\begin{aligned} & \text{Tr}(\mathbf{C}_{,\alpha} \mathbf{\Pi}^T \mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1} \mathbf{\Pi}) \\ &= \frac{1}{2} \text{Cov}(\mathbf{x}^T \mathbf{\Pi}^T \mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{x}, \mathbf{x}^T \mathbf{\Pi}^T \mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{x}) \end{aligned} \quad (53)$$

where \mathbf{x} is normally distributed with covariance 1. For Gaussian \mathbf{x} , the covariance can be approximated as the $2\nu(k)^{-1} P_\alpha^2 \delta_{\alpha,\beta}$, where P_α is the mean power spectrum (in this case $\langle \mathbf{x}^T \mathbf{\Pi}^T \mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{x} \rangle$) and $\nu(k)$ is the number of modes in the survey volume. A similar expression holds for the denominator of the transfer function. The numerator can be estimated in Monte Carlo with lower noise by removing the residual foreground variance as in Equation 52. The $\nu(k)$ pre-factors drop out in the ratio, and the transfer function can then be expressed as the average of the ratios

$$T_\alpha^F \approx \left\langle \frac{[\mathbf{\Pi}_{\text{s+f}}(\mathbf{x} + \mathbf{x}_{\text{sim}}) - \mathbf{\Pi}_f \mathbf{x}]^T \mathbf{C}^{-1} \mathbf{Q}_\alpha \mathbf{x}_{\text{sim}}}{\mathbf{x}_{\text{sim}}^T \mathbf{C}^{-1} \mathbf{Q}_\alpha \mathbf{x}_{\text{sim}}} \right\rangle^2. \quad (54)$$

Equation 54 provides a convenient procedure for estimating signal attenuation due to cleaning empirically determined spectral modes.

Figure 5 shows the transfer function for the GBT-wide data (including the beam). At high k_\perp , signal is lost to large beam size. The signal attenuation at low k_\perp and low k_\parallel is due to spectral foreground cleaning. Signal attenuation toward low k_\parallel can be explained as “direct” loss of removing functions along the line of sight that have overlap with the signal’s spectral variation (Section 4.2). The structure at $k_\parallel = 0.07 \text{ hMpc}^{-1}$ originates from additional foreground degrees of freedom around that scale that needed to be nulled. The signal loss at low k_\perp is less intuitive and is due to the spurious spatial correlation of signal and foreground on large angular scales. Each foreground spectral mode \mathbf{u}_i^f discovered in the survey is associated with a spatial mode through the singular value decomposition (SVD) $\mathbf{X} = \sqrt{N_\theta} \mathbf{U} \mathbf{\Lambda}^{1/2} \mathbf{V}^T$ (see Appendix A and Nityananda (2010)). The spectral functions nulled in the line of sight are associated with spatial modes of the signal. The largest foregrounds are spatially smooth and share some spurious correlation with the small number of spatial signals at low k_\perp . Nulling the spectral variation associated with these spatial modes erases the large spatial scales. Signal loss toward low k_\perp is an artifact of nulling spectral modes determined from the map itself.

5.4. Aggressiveness of Foreground Cleaning

For residual foregrounds to be negligible, (1) the foreground must inherently have few spectral degrees of freedom compared to the signal and (2) instrument response must be well constrained and mix bright foregrounds into a limited number of new modes. Some instrument responses like calibration stability will inevitably produce a long tail of new modes where each line of sight observes a slightly different continuum. This produces high-rank residuals that cannot be easily estimated or subtracted.

Until proven otherwise, interpretation of intensity mapping power spectra should account for the additive bias from residual/unweighted foregrounds from misspecification or incomplete knowledge of the foreground covariance. The goal of good instrument design and calibration is to push the amplitude of residual modes well below the thermal noise in the




Figure 5. Transfer function of signal attenuation for 30 foreground modes removed in the GBT-wide field, including the effect of the beam. When contaminated modes are determined from the data themselves, it is essential to include foregrounds in simulations of the signal attenuation. This can be viewed two ways: either (1) the signal perturbs the foreground modes or (2) spurious correlations between signal and foreground need to be included. Especially on k -modes that approach the size of the survey region, there are few modes available and most of the signal is lost. To reach the lowest upper bound of the signal auto-power, foreground cleaning also removed much of the signal. The beam rolls off response at high k_{\perp} , and the signal attenuation at low k_{\perp} and k_{\parallel} is due to the spectral foreground removal.

maps, so that the biases are at the level of the power spectral errors. (For some goals with particular $P(k)$ shapes like the baryon acoustic oscillation, requirements on the additive bias may be weaker.)

Cross-correlation with a spectroscopic galaxy catalog has the advantage that residual foregrounds will boost errors rather than producing an additive bias. Galaxy densities and line intensity fields will not be perfectly correlated, so cross-correlation provides a lower bound on the fluctuation power in the line survey. Combined with the auto-power of the line survey (which is an upper limit due to its additive bias), an intensity survey can provide an indirect inference on the range of true fluctuation power in the atomic line (Switzer et al. 2013).

Foreground cleaning requirements vary between auto-power surveys of intensity maps and cross-powers with side surveys. For too few modes removed, the cross-power errors are large because of the map variance from foregrounds. By removing more than the optimal number of modes, the error bars increase again as more signal is lost (Masui et al. 2013). The optimal cleaning is achieved for the lowest errors on the cross-power, a well-defined quantity. Similarly, the additive bias on the auto-power of the intensity mapping survey drops as more modes are removed, but errors again increase. The upper bound on the line fluctuation intensity is the band power plus the error. One can choose a number of modes to remove that gives the lowest upper bound on the auto-power amplitude. Generally, the foreground cleaning requirements for the cross-power with a side survey are much less stringent than the auto-power of the intensity survey. Each Fourier mode of the cross-power has a different realization of the residual foregrounds crossed with the galaxy survey. In a survey with many k -modes contributing to a band power, the cross-powers can average down over these correlations to produce small errors. In contrast, each mode of the auto-power of the intensity map with itself is a sample of approximately the same power, which does not average down.

Figure 6 shows a high-level summary of the plausibility of

detecting the auto-power in line intensity surveys. It is clear that one wants (1) a survey that has many signal degrees of freedom in the k -modes of interest and (2) a shallow eigenvalue spectrum of the foregrounds. These two goals can be contradictory. The first goal suggests coverage of wide areas of the sky to suppress spurious correlations between signal and foreground. However, in surveying large areas, an instrument may observe a wider range of foreground spectral modes or be prone to instrumental effects that are harder to control across large areas and time. These effects generally boost the number of modes that need to be estimated. Possible examples of these effects are spatial variations in spectral index, variations in bandpass calibration, or exposure to a wider range of rotation measures through polarization to intensity leakage.

CMB B-modes have had a vigorous history of studies of instrumental systematics (e.g. Hu et al. (2003)). Similar studies should be undertaken for intensity mapping surveys. These would characterize the impact of calibration stability, beam response, or other instrumental response to foregrounds. This requires detailed simulations that are beyond the scope of this methods paper. Foreground emission and variations in spectral index may contribute a handful of spectral degrees of freedom, but instrumental effects have the potential to mix these into many more bright degrees of freedom that are not well known in advance. In future studies, the eigenvalue spectrum of input astronomical foregrounds can be compared to the eigenvalue spectrum as observed by a simulated instrument. A well-designed survey will control the number of modes induced by the instrument response, targeting the rapidly falling green contours in Figure 6.

6. ASSEMBLING THE FINAL PRODUCT

Previous sections defined the core aspects of the estimator and foreground cleaning. In practice, the estimator can be made more robust to temporally variable noise by forming cross-powers between subseasons. The 2D band powers derived above are expected to have low signal-to-noise ratio in the first generation of experiments. They are also difficult to interpret. We describe a procedure to optimally bin onto 1D powers and develop a covariance model and decorrelation of those powers for convenient display. The binning weights from 2D to 1D provide insight into the information content of the intensity survey.

6.1. The Subseason Cross-power

Thermal noise is uncorrelated between subseasons of the observations to an excellent approximation. In addition, some forms of contamination such as time-varying RFI, or calibration instability (and its induced foreground residuals), may be largely uncorrelated across times. The cross-powers between N_s split subseasons of the data $\mathbf{x}_A \dots \mathbf{x}_{N_s}$ therefore have no additive noise bias from these terms (Tristram et al. 2005; Switzer et al. 2013; Dillon et al. 2014). A subseason cross-power estimator is formally nonoptimal, but does follow from the form of the optimal quadratic estimator. Appendix B shows the choices that lead to the estimator

$$\hat{q}_{\alpha}^{A \times B} \propto (\mathbf{N}_A^{-1} \mathbf{x}_A)^T \mathbf{C}_{\alpha} (\mathbf{N}_B^{-1} \mathbf{x}_B), \quad (55)$$

where \mathbf{N}_A and \mathbf{N}_B are the nonsignal noise covariance in two maps. The data are weighted by their respective covariances for the two subseasons. Following the form of Equation 23 to




Figure 6. Intensity mapping foreground regimes. The signal power at a fixed k band power is a constant (black) normalized to 1. The eigenvalue spectrum of foregrounds is approximated as a power law, with a steep (green) and shallow (red) decline from 10^4 times the signal (in power, not map space), as more modes are removed. Errors on the band power (blue) increase as more signal is removed by foreground cleaning, according to Equation 48. The noise level is set to achieve a 5σ detection of the signal at zero modes removed. Residual foregrounds are also boosted by the signal loss transfer function. Dashed red and green show the decline of the intrinsic foreground variance as more modes are removed. The solid lines show the residual foreground variance with transfer function applied to keep signal constant. A successful experiment will have foregrounds below noise and noise below signal. The left panel shows an experiment that has 80 spatial degrees of freedom for this k mode. This experiment will only be successful if the foreground eigenvalue spectrum falls rapidly (green). Note that for the shallow spectrum, the contamination at its minimum is hundreds of times the signal and larger than the noise. The right panel shows an experiment that has 250 degrees of freedom for the k mode. In this case, there are sufficient modes to secure a detection even for the shallow foreground eigenvalue spectrum. This does not include the impact of noise on contaminated mode determination, which generally limits the maximum number of modes that can be removed, or alternately how far the foreground can be suppressed relative to thermal noise.

avoid contaminated modes,

$$\hat{q}_\alpha^{A \times B} \propto (\mathbf{W}_A \mathbf{\Pi}_A \mathbf{x}_A)^T \mathbf{C}_{\alpha} (\mathbf{W}_B \mathbf{\Pi}_B \mathbf{x}_B). \quad (56)$$

This expression continues to assume that after projecting the contaminated modes through $\mathbf{\Pi}_A \mathbf{x}_A$, the remaining covariance is dominated by thermal noise, which simply follows the coverage map of the experiment in the two subseasons and can be implemented as diagonal inverse covariance weights \mathbf{W}_A and \mathbf{W}_B . Take the projections to remove spectral modes along each line of sight as in Equation 29,

$$\mathbf{\Pi}_A \mathbf{x}_A = \text{vec}(\mathbf{\Pi}_A^\nu \mathbf{X}_A) = \text{vec}((\mathbf{1} - \mathbf{U}_{A,f} \mathbf{S} \mathbf{U}_{A,f}^T), \mathbf{X}_A) \quad (57)$$

where now the contaminated modes $\mathbf{U}_{A,f}$ can be particular to the subseason. (If the instrument noise varies with frequency, that weight should appear here as in Equation 27. We neglect it to keep a simpler equation.)

We can also develop foreground modes that are best tuned to the respective data splits by finding modes of the cross-covariance of the two subseason maps as

$$\hat{\mathbf{C}}_{\nu,\nu'} = \mathbf{X}_A \mathbf{X}_B^T = \mathbf{U}_{A,f} \mathbf{\Lambda} \mathbf{U}_{B,f}^T, \quad (58)$$

where the final operation is the SVD of the cross-variance. The covariance estimate can be improved by using the weighted maps $\mathbf{W}_A \mathbf{x}_A$ (Switzer et al. 2013). However, in doing this it is important to force \mathbf{W}_A to be a separable expression as $\mathbf{W}_A^\nu \otimes \mathbf{W}_A^\theta$ by averaging over the frequency and spatial directions. Otherwise, the weighting operation would increase the rank (e.g. nonseparable \mathbf{W} applied to $\mathbf{X} = \mathbf{u} \mathbf{v}^T$ is no longer generally rank-1).

Most instruments have a chromatic beam. Without accounting for this, the modal structure will include variance introduced by this frequency response. In GBT, this was the second-largest mode (Figure 4) and was treated by convolving the maps to a common resolution. In general, any well-modeled aspects of the instrument should be treated in the map to reduce the burden of discovering those modes in the

data.

The estimate \hat{q}_α is the average over all cross-power pairs

$$\hat{q}_\alpha = \frac{1}{N_s(N_s - 1)} \sum_{i,j,i \neq j} \hat{q}_\alpha^{i \times j}. \quad (59)$$

We have assumed that all pairs have similar statistical power. If thermal noise dominates, this can be arranged by forming all of the subseason split maps from approximately the same integration times and areas. If this is not possible, then the cross-powers should be more optimally weighted.

6.2. Projecting to 1D Powers

The optimal estimator structure developed in Section 3.4 binned from 3D Fourier space to the 2D band powers assuming that the noise was isotropic across constant $k_\perp = \sqrt{k_x^2 + k_y^2}$ under the action of \mathbf{C}^{-1} . For the first generation of intensity mapping experiments, there is likely to be insufficient signal-to-noise ratio on each 2D band power. Binned 1D power $\hat{P}(k)$ for a band of $\sqrt{k_\perp^2 + k_\parallel^2} \in (k, k + \Delta k)$ contains most of the cosmological information other than redshift-space distortions and so provides a convenient summary.

Noise is strongly anisotropic across the k_\perp and k_\parallel ring that contributes to some k . This is because the number of modes increases quadratically in k_\perp and linearly in k_\parallel , the transfer functions will vary across the ring in 2D k -space, and the noise itself may vary across the ring. These factors also make it difficult to interpret the 2D power spectrum. In going from 2D to 1D powers, we can apply an inverse-covariance weight to maximize the 1D signal-to-noise ratio. Let α continue to refer to 2D band powers $k_{\perp,\alpha}, k_{\parallel,\alpha}$ and R_k be the ring of k_α values that contribute to the k band power. The weighted sum is

$$\hat{P}(k) = \frac{\sum_{\alpha \in R_k} I_{\alpha \in R_k} W_\alpha \hat{p}_\alpha}{\sum_{\alpha \in R_k} I_{\alpha \in R_k} W_\alpha}. \quad (60)$$

Figure 7 shows the indicator $I_{\alpha \in R_k}$ as a set of colored bands

for each k_α bin. Again the choice of W_α is a weighting related to optimality rather than bias so long as W_α is based on the variance rather than the mean value of the bins. In contrast, an analysis should not use W_α that does additional foreground suppression by masking areas of the 2D power that have apparent residual foregrounds in the mean band power, as this would produce bias.

In the language of subseason data splits, $\hat{p}_\alpha^{A \times A}$ is an auto-power, and it describes the variance of thermal noise, cosmological signal, and residual foregrounds. In practice for GBT-wide, this was dominated per k_α bandpower cell by thermal noise. Assuming Gaussian errors from the dominant thermal noise, the inverse covariance weight per pixel is

$$W_\alpha = \frac{N_\alpha T_\alpha^2}{\hat{p}_{\text{auto},\alpha}^2}, \quad (61)$$

where N_α is the number of 3D Fourier cells contributing to the 2D power in band power α , T_α^2 is the total transfer function that applies to the band power α , and $\hat{p}_{\text{auto},\alpha}$ is estimated from the 2D auto-powers. Figure 7 shows W_α for the GBT-wide survey. Prior to these measurements, the theoretical expectation was that foreground spectra vary slowly and that primarily information at low k_\parallel is lost. In the case of foreground modes determined from the data themselves, large angular scales have more spurious correlation with foregrounds, and considerable k_\perp is also lost, resulting in deweighting of low k_\perp . Interband correlations are estimated in the next section using simulations, but are not exploited in the 2D to 1D weighting.

6.3. Errors and Decorrelation

A simple estimate of the band power errors can be derived from the variance of cross-power pairs of subsets of the data. In the case of a survey split into maps A , B , C , and D , the cross-powers $A \times B$, $A \times C$, $A \times D$, $B \times C$, $B \times D$, $C \times D$ form six unique samples of the instrument's thermal noise. The subseason maps A , B , C and D all share the same underlying signal and residual foregrounds, so the variance of the crossed pairs does not reflect the sample variance. It is also a poor estimate of the errors if there are a limited number of crossed pairs (six in the example here).

If the cleaned maps are dominated by Gaussian fluctuations such as from thermal noise, the full bandpower covariance can be determined from the measured auto-power spectrum between common sections (e.g. $A \times A$), the power spectrum across sections, and the survey geometry (see, e.g. Das et al. (2011) for an application in CMB analysis). The validity of Gaussian errors depends on both the signal and residual foregrounds. If a survey resolves sufficient non-Gaussian cosmological signal, then methods such as Harnois-Déraps & Pen (2012) should be considered to properly describe the errors. To date, no intensity maps have reached this high cosmological signal-to-noise regime. Appendix C calculates the full band power covariance $\text{cov}(\hat{P}, \hat{P})$ using a hybrid of Monte Carlo simulations for the off-diagonal structure and Gaussian errors for the amplitudes.

With the optimal estimator, both the final covariance and the window function are the Fisher matrix. Standard discussion of decorrelation (Hamilton & Tegmark 2000) for optimal estimators freely moves between undoing the effect of the window function through \mathbf{F}^{-1} and diagonalizing the final band power covariance through $\mathbf{F}^{-1/2}$. For the suboptimal estimators here, the variance $2\text{Tr}(\mathbf{CQ}_\alpha \mathbf{CQ}_\beta)$ is dif-

ferent from the windowing matrix $\text{Tr}(\mathbf{C}_\beta \mathbf{Q}_\alpha)$. The band powers \hat{p}_α so far have just been scalar normalizations times the pseudo-powers \hat{q}_α rather than a linear combination that decorrelates the band powers.

With the full band power covariance model in hand from Appendix C, we can repeat the classic decorrelation choice (Hamilton & Tegmark 2000) of $\mathbf{F}^{-1/2}$ by taking $\text{cov}(\hat{P}, \hat{P})_{\text{data}}^{-1/2}$ to multiply by the band powers (normalized so that the weights on the band powers sum to 1). So long as decorrelation multiplies by an invertible matrix, no information is lost and the choice is purely one of display, which generally benefits from uncorrelated errors. Switzer et al. (2013) and Masui et al. (2013) use this pipeline for two GBT intensity mapping surveys and describe results and the interpretation of 1D band powers subject to additive bias.

7. DISCUSSION

Intensity mapping experiments have the potential to map cosmological volumes with resolution and sensitivity requirements that are modest compared to direct spectroscopic surveys of objects. In addition to atomic or molecular line radiation, these surveys generally receive continuum radiation that can be orders of magnitude brighter. We have developed a quadratic estimator that combines some aspects of both galaxy and CMB surveys, but also accommodates methods of down-weighting bright continuum emission. A fully optimal estimator requires a model of the covariance of contamination, which we argue is not well known prior to an experiment. In the example of GBT data, the spectral structure of contaminant modes was related primarily to the instrument response rather than intrinsic spectra. The instrumental response is residual in the sense that considerable new effort was put into calibration, in addition to the heritage of a well-established instrument.

We develop an estimator that (1) takes cross-powers of subseasons to avoid bias from temporally-variable noise (2) estimates a foreground covariance model from the data itself through a reduction in dimensionality, (3) accounts for the impact of spurious correlations between signal and foreground, and (4) derives the final 1D power and its errors. Transfer functions provide a convenient way to calibrate the estimator's output to the signal input and can be estimated efficiently using Monte Carlo simulations. Spurious correlations of signal and foreground result in an average anticorrelation of signal and residual foregrounds. Simulations for the transfer function must include foregrounds to properly account for this effect.

Figure 6 argues that for an intensity mapping experiment to be successful, it must control the eigenvalue spectrum of foregrounds and observe a large enough area that spurious correlation between signal and foreground can average down. While the intrinsic foregrounds may only have a handful of degrees of freedom, variations in instrumental response have the potential to mix those spectral modes into a larger number of new modes and a shallower eigenvalue spectrum. In particular, variable spectral calibration contributes some level of full-rank covariance (each line of sight responds differently to bright emission), even with rank-1 input contamination. The eigenvalue spectrum of the ν, ν' covariance of the maps is the central metric for the quality of the calibration or mapping procedure. Because of instrumental effects, contaminant modes are not necessarily smooth and so generally poorly described by smooth functions. The salient aspect here is not




Figure 7. Left: binning from 2D to 1D powers occurs in bands along constant $k = \sqrt{k_{\perp}^2 + k_{\parallel}^2}$ for 15 bands between $k = 0.1 \text{ hMpc}^{-1}$ and 1 hMpc^{-1} . Constant colors show the 2D k -cells contributing to a 1D band power. The bins are constant in $\log(k)$. Right: the noise weight from Equation 61 for the GBT-wide survey. The weight bins per 2D bandpower are normalized to sum to 1 across the 1D k -bands in the left panel. The colored bands in the left panel guide the eye for 2D k bins that contribute. At low $k \approx 0.1 \text{ hMpc}^{-1}$, the weight favors slowly varying spectral modes and more rapidly varying spatial modes (the vertical part of the iso- k contour). This reflects the fact that the foreground cleaning destroys spatial modes with $k_{\perp} < 0.1 \text{ hMpc}^{-1}$ because the bright foreground modes are associated with spatially smooth structure in the map. In contrast, long-wave spectral modes are less penalized because the contaminated modes have high-frequency spectral structure (only the largest mean synchrotron mode is spectrally smooth). In contrast, by $k \approx 1 \text{ hMpc}^{-1}$ almost all information comes from rapidly varying spectral modes (the horizontal part of the iso- k contour) because the beam has destroyed essentially all of the information at $k_{\perp} > 0.4 \text{ hMpc}^{-1}$. For intermediate scales, information at a wider range of modest k_{\perp} and k_{\parallel} is weighted more heavily. The noisy structure of the weights is driven by variations in the auto-power in the denominator of Equation 61.

the spectral smoothness but rather that the signal can fluctuate in many more ways than the instrument's response to bright foregrounds.

Figure 6 also demonstrates one of the challenges of reaching convincing detection using intensity mapping data alone, when no cross-power corroboration is possible. An experiment only has access to the band power estimate of signal plus foreground, which formally represents an upper bound on signal. This total bandpower falls as foregrounds are more aggressively cleaned. If signal dominates, the bandpower will reach a plateau where errors increase, but the amplitude does not diminish as the cleaning pushes to down-weight more foreground structure. However, a shallower plateau could also result from the fact that residual foreground variance is boosted after accounting for the transfer function (see the solid red curve, left panel of Figure 6). The onus is to argue (1) that cosmological signal power is detected and is stable to efforts to clean additional foregrounds and (2) that residual foregrounds and the signal transfer function do not conspire to appear as a stable signal band power. Additionally, there may be features in the power spectrum such as the BAO feature, or redshift-space distortions (in the 2D spectrum), that

support the interpretation of cosmological signal.

Intensity mapping shares some parallels with CMB B-mode searches, where instruments must be designed to prevent mixing between bright contaminants and the signal, and foreground cleaning is a central strategy. The same language and metrics that have been developed for beam systematics in B-mode searches would be fruitfully carried over to intensity mapping. All of the lines and redshift ranges of interest have differences in experimental methodology, but the eigenvalue spectrum provides a common reference for developing instrumental requirements. Beyond astrophysical foregrounds, planning of future intensity mapping experiments should also include high-fidelity simulations of the instrument to determine requirements for the accuracy and stability of the spectral calibration.

E.S. acknowledges support as a CITA fellow, where much of this work was conducted. We thank John Ford, Anish Roshi, and the rest of the GBT staff for their support, especially in understanding instrument response. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by the Canada Foundation for Innovation.

APPENDIX

IMPACT OF FOREGROUND-SIGNAL COUPLING IN BLIND CLEANING

In this appendix, we find the expectation value of the correlation between the residual foregrounds in the cleaned map and the input signal. Section 5.2 argues that the cleaned maps contain $\tilde{\mathbf{X}}_s = -\Delta \mathbf{S} \mathbf{U}_f^T \mathbf{X}_f - \mathbf{U}_f \mathbf{S} \Delta^T \mathbf{X}_f$. The term Δ describes how pure foreground spectral modes are influenced by the signal. In this appendix, we argue that $\tilde{\mathbf{X}}_s$ is anticorrelated with the signal, on average. At first order in a perturbing signal, $\Delta = \mathbf{U}_f \mathbf{H}$ where the matrix elements are

$$\mathbf{H}_{ij} = N_{\theta}^{-1} \frac{(\mathbf{u}_i^f)^T (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T) \mathbf{u}_j^f}{\lambda_j^f - \lambda_i^f}. \quad (\text{A1})$$

Here $\mathbf{H} = -\mathbf{H}^T$, is skew-symmetric because of the denominator of the perturbation element, while the numerator is symmetric by construction. The correlation between residual foregrounds in $\tilde{\mathbf{X}}_s$ and signal due to spurious correlations is

$$\xi_{\text{spur}} = N_\theta^{-1} \langle \text{Tr}(\tilde{\mathbf{X}}_s \mathbf{X}_s^T) \rangle = -N_\theta^{-1} \langle \text{Tr}(\mathbf{U}_f \mathbf{H} \mathbf{S} \mathbf{U}_f^T \mathbf{X}_f \mathbf{X}_s^T + \mathbf{U}_f \mathbf{S} \mathbf{H}^T \mathbf{U}_f^T \mathbf{X}_f \mathbf{X}_s^T) \rangle \quad (\text{A2})$$

$$= -N_\theta^{-1} \langle \text{Tr}(\mathbf{H} \mathbf{S} \mathbf{U}_f^T \mathbf{X}_f \mathbf{X}_s^T \mathbf{U}_f + \mathbf{S} \mathbf{H}^T \mathbf{U}_f^T \mathbf{X}_f \mathbf{X}_s^T \mathbf{U}_f) \rangle \quad (\text{A3})$$

$$= -N_\theta^{-1} \langle \text{Tr}(\mathbf{H} \mathbf{S} \mathbf{U}_f^T (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T) \mathbf{U}_f) \rangle, \quad (\text{A4})$$

where the first line uses $\Delta = \mathbf{U}_f \mathbf{H}$, the second line uses the cyclic property of the trace, and the third line uses the fact that a matrix and its transpose have the same trace, and the symmetry $\mathbf{S}^T = \mathbf{S}$. Let $\Sigma = \mathbf{U}_f^T (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T) \mathbf{U}_f$ and note that $\mathbf{H}|_{ij} = N_\theta^{-1} \Sigma|_{ij} \cdot (\lambda_j^f - \lambda_i^f)^{-1}$. Here Σ is a symmetric cross-variance of the signal and foreground in the basis of the foreground modes. We can split all matrices in the foreground mode basis into cut modes and uncut modes. The matrix \mathbf{S} is 1 along the diagonal for cut modes and zero elsewhere. Order the cuts so that they all reside in the upper left submatrix as

$$\mathbf{S} = \begin{pmatrix} \mathbf{1}_{cc} & \mathbf{0}_{cu} \\ \mathbf{0}_{uc} & \mathbf{0}_{uu} \end{pmatrix} \quad \Sigma = \begin{pmatrix} \Sigma_{cc} & \Sigma_{cu} \\ \Sigma_{uc} & \Sigma_{uu} \end{pmatrix}, \quad (\text{A5})$$

where c and u denote modes that are cut vs. uncut, and analogously for \mathbf{H} . The trace evaluates to

$$\xi_{\text{spur}} = -N_\theta^{-1} \langle \text{Tr}(\mathbf{H}_{cc} \Sigma_{cc}) + \text{Tr}(\mathbf{H}_{uc} \Sigma_{cu}) \rangle. \quad (\text{A6})$$

The first term is the trace of the product of skew-symmetric and symmetric matrices, so it is zero. Evaluating the second term using the symmetry of Σ produces

$$\xi_{\text{spur}} = -N_\theta^{-2} \left\langle \sum_{\substack{i \in \text{cuts} \\ j \notin \text{cuts}}} \frac{[(\mathbf{u}_i^f)^T (\mathbf{X}_f \mathbf{X}_s^T + \mathbf{X}_s \mathbf{X}_f^T) \mathbf{u}_j^f]^2}{\lambda_i^f - \lambda_j^f} \right\rangle, \quad (\text{A7})$$

taking the transpose of \mathbf{H} and reversing the denominator to preserve sign. The elements of the spurious correlation $\Sigma|_{ij}$ have zero mean (over signal realizations), but the correlation between signal and cleaning residuals in ξ_{spur} appears quadratically inside the average over signal realizations, resulting in a net bias.

Take the SVD of \mathbf{X}_f , $\mathbf{X}_f = \sqrt{N_\theta} \mathbf{U}_f \mathbf{\Lambda}_f^{1/2} \mathbf{V}_f^T$ so that

$$\xi_{\text{spur}} = -N_\theta^{-1} \left\langle \sum_{\substack{i \in \text{cuts} \\ j \notin \text{cuts}}} \frac{[\sqrt{\lambda_i^f} (\mathbf{v}_i^f)^T \mathbf{X}_s^T \mathbf{u}_j^f + \sqrt{\lambda_j^f} (\mathbf{u}_i^f)^T \mathbf{X}_s \mathbf{v}_j^f]^2}{\lambda_i^f - \lambda_j^f} \right\rangle \quad (\text{A8})$$

By the construction of the filter, $\lambda_i^f \gg \lambda_j^f$ for $i \in \text{cuts}$ and $j \notin \text{cuts}$ because the cuts remove the highest variance foreground modes. In this limit,

$$\xi_{\text{spur}} = -N_\theta^{-1} \left\langle \sum_{\substack{i \in \text{cuts} \\ j \notin \text{cuts}}} [(\mathbf{v}_i^f)^T \mathbf{X}_s^T \mathbf{u}_j^f]^2 \right\rangle. \quad (\text{A9})$$

The amplitudes of foregrounds drop out, and the matrix element $(\mathbf{v}_i^f)^T \mathbf{X}_s^T \mathbf{u}_j^f$ is the overlap of the signal with the subtracted foreground spatial modes and unsubtracted spectral modes. If the subtracted modes have a smooth spatial distribution, spurious correlations will wipe out signal at low k_\perp . Using the SVD of the signal $\mathbf{X}_s = \sum_n \sqrt{\lambda_n^s} N_\theta \mathbf{u}_i^s (\mathbf{v}_i^s)^T$,

$$\xi_{\text{spur}} = - \left\langle \sum_{\substack{i \in \text{cuts} \\ j \notin \text{cuts}}} \left[\sum_n \sqrt{\lambda_n^s} (\mathbf{v}_i^f)^T \mathbf{v}_n^s (\mathbf{u}_n^s)^T \mathbf{u}_j^f \right]^2 \right\rangle. \quad (\text{A10})$$

To expand the squared sum on n , note that cross terms with $n \neq n'$ will average to zero in the $\langle \rangle$ over signal. Recall that both the spatial and spectral modes are normalized so that $\mathbf{u}^T \mathbf{u} = 1$ and $\mathbf{v}^T \mathbf{v} = 1$. To get a rule of thumb, let $N_{\text{res},\nu}$ be the effective number of spectral degrees of freedom of the signal fluctuation. Then each inner product $[(\mathbf{u}_k^s)^T \mathbf{u}_j^f]^2 \approx 1/N_{\text{res},\nu}$ and the spatial inner products-squared scale as $\approx 1/N_{\text{res},\theta}$. Let the sum on $i \in \text{cuts}$ be over N_m cut modes. If there are $N_{\text{res},\nu}$ spectral modes available in the survey, then the number of uncut modes in the sum $j \notin \text{cuts}$ is over $N_{\text{res},\nu} - N_m$. Using the sum of the signal eigenvectors from Equation 31, the scaling of ξ_{spur} gives the rule of thumb

$$\xi_{\text{spur}} \approx -\xi_s \frac{N_m (N_{\text{res},\nu} - N_m)}{N_{\text{res},\nu} N_{\text{res},\theta}}. \quad (\text{A11})$$

Recall that $\xi_{\text{clean}} = \xi_s + \xi_{\text{direct}} + \xi_{\text{spur}}$, or

$$\xi_{\text{clean}} = \xi_s \left[1 - \frac{N_m}{N_{\text{res},\nu}} - \frac{N_m (N_{\text{res},\nu} - N_m)}{N_{\text{res},\nu} N_{\text{res},\theta}} \right] = \left(1 - \frac{N_m}{N_{\text{res},\nu}} \right) \left(1 - \frac{N_m}{N_{\text{res},\theta}} \right) \xi_s. \quad (\text{A12})$$

In interpreting this rule of thumb, it is useful to think about the survey volume filtered onto particular scales for a given band power. The terms $N_{\text{res},\nu}$ and $N_{\text{res},\theta}$ are essentially unrelated to the number of frequency bins N_ν and spatial pixels N_θ in the survey. Either N_ν or N_θ could be made arbitrarily large through mapping with finer pixels or a larger number of spectral channels. Instead, the relevant quantity is the number of spatial and spectral degrees of freedom that the signal in the given 2D band power can explore. If there can be many signal realizations on a given scale, then the spurious correlation with the foregrounds averages down better.

RELATION OF THE CROSS-POWER TO THE OPTIMAL ESTIMATOR

Split the season into two maps $\mathbf{x}^T = (\mathbf{x}_A, \mathbf{x}_B)^T$ and let the covariance be

$$\mathbf{C} = \begin{pmatrix} \mathbf{S} + \mathbf{N}_A & \mathbf{S}_\times \\ \mathbf{S}_\times & \mathbf{S} + \mathbf{N}_B \end{pmatrix}, \quad (\text{B1})$$

where $\mathbf{S} = \mathbf{S}_\times = \sum_\alpha p_\alpha \mathbf{S}_{,\alpha}$ is the signal covariance and $\mathbf{N}_A, \mathbf{N}_B$ are the noise covariance in the two maps. We assume that the noise covariance contains only thermal noise, which is uncorrelated between subseasons. Any signal on the sky (including residual foregrounds) is correlated between subseasons, and we absorb it in the signal covariance. The optimal estimator remains

$$\hat{q}_\alpha \propto \mathbf{x}^T \mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1} \mathbf{x}. \quad (\text{B2})$$

This generically involves combinations of the data like $\mathbf{x}_A^T \mathbf{Q} \mathbf{x}_A$ (auto-powers) and $\mathbf{x}_A^T \mathbf{Q} \mathbf{x}_B$ (cross-powers). To avoid noise bias, we would like to avoid terms like $\mathbf{x}_A^T \mathbf{Q} \mathbf{x}_A$. This is done by 1) neglecting the blocks along the diagonal of $\mathbf{C}_{,\alpha}$, and 2) by neglecting the signal covariance contribution to \mathbf{C} . These choices are

$$\mathbf{C}_{,\alpha} \Big|_{\text{cross}} = \begin{pmatrix} 0 & \mathbf{S}_{,\alpha} \\ \mathbf{S}_{,\alpha} & 0 \end{pmatrix}, \quad \mathbf{C} \Big|_{\text{cross}} = \begin{pmatrix} \mathbf{N}_A & 0 \\ 0 & \mathbf{N}_B \end{pmatrix} \Rightarrow b_\alpha = \text{Tr}(\mathbf{C}_{,\alpha} \mathbf{C}) = 0. \quad (\text{B3})$$

Putting these factors together, the crossed estimator is

$$\hat{q}_\alpha \propto (\mathbf{N}_A^{-1} \mathbf{x}_A)^T \mathbf{S}_{,\alpha} (\mathbf{N}_B^{-1} \mathbf{x}_B)^T. \quad (\text{B4})$$

Formally, this cross-power is suboptimal because it neglects signal correlations in the weighting, and it neglects signal information in the auto-power.

GAUSSIAN ERRORS

This derivation follows Das et al. (2011) except that we do not form explicit map differences to estimate thermal noise. The covariance of an estimator $\hat{P}_{i \times j}$ across Gaussian fields i, j is

$$\text{cov}(\hat{P}_{i \times j}, \hat{P}_{k \times l}) = \frac{1}{\nu(k)} [\langle P_{i \times k} \rangle \langle P_{j \times l} \rangle + \langle P_{i \times l} \rangle \langle P_{k \times j} \rangle], \quad (\text{C1})$$

where $\nu(k)$ is the effective number of modes that enter the average for the band power. Model the power spectra as $\langle P_{i \times j} \rangle = P_{\text{auto}}$ for $i = j$ and $\langle P_{i \times j} \rangle = P_\times$ for $i \neq j$. For simplicity, we will assume that each map section has approximately the same statistical properties so that the cross-powers are represented by P_\times (e.g. $A \times B$) and the auto-powers are represented by P_{auto} (e.g. $A \times A$) to a good approximation. However, in surveys where map sections have different integration times or noise properties, these expressions should be expanded to break out the noise properties of the different sub-surveys. Note that P_{auto} includes both thermal noise and sky variance, and P_\times includes any sky variance (including residual foregrounds).

The covariance in Equation C1 of several data combinations is

$$\text{no sec. in common} \quad \frac{2}{\nu(k)} P_\times^2 \quad (\text{C2})$$

$$\text{one sec. in common} \quad \frac{1}{\nu(k)} [P_\times^2 + P_\times P_{\text{auto}}] \quad (\text{C3})$$

$$\text{two sec. in common} \quad \frac{1}{\nu(k)} [P_\times^2 + P_{\text{auto}}^2]. \quad (\text{C4})$$

An example of the first case would be AB, CD , the second case would be AB, AC and the third case, AB, AB . The total covariance of the estimated power spectrum is the sum of the covariance cases above with appropriate multiplicities

$$\begin{aligned} \text{cov}(\hat{P}(k), \hat{P}(k)) &= \frac{1}{N_s(N_s - 1)} \left[(N_s - 2)(N_s - 3) \frac{2}{\nu(k)} P_\times^2 + 4(N_s - 2) \frac{1}{\nu(k)} (P_\times^2 + P_\times P_{\text{auto}}) + \frac{2}{\nu(k)} (P_\times^2 + P_{\text{auto}}^2) \right] \\ &= \frac{1}{\nu(k) N_s(N_s - 1)} [2(N_s^2 - 3N_s + 3) P_\times^2 + 4(N_s - 2) P_\times P_{\text{auto}} + 2P_{\text{auto}}^2]. \end{aligned} \quad (\text{C5})$$

We can cast this in a more familiar form by letting $P_\times = P_{\text{sig}}$ be the ‘‘signal’’ and $P_{\text{auto}} = P_{\text{sig}} + P_n$ ‘‘signal plus noise.’’ Separating the signal and noise powers from P_{auto} one has

$$\nu(k) \text{cov}(\hat{P}(k), \hat{P}(k)) = 2P_{\text{sig}}^2 + 4 \frac{P_{\text{sig}} P_n}{N_s} + 2 \frac{P_n^2}{N_s(N_s - 1)}. \quad (\text{C6})$$

These Gaussian errors determine the diagonal of the bandpower covariance and require an estimate of the number of modes $\nu(k)$. In practice, restrictions of the survey volume due to spatial/spectral weighting and masking produce a complex covariance structure and effective number of modes. We advocate a hybrid approach where the bandpower covariance diagonal is estimated through Gaussian errors using the auto- and cross-powers of the data (and so represents sample variance and thermal noise). This is then used to calibrate a full covariance matrix determined by Monte Carlo of the complete pipeline.

Let $\mathbf{C}_{\text{sim}} = \text{cov}(\hat{P}(k)_{\text{sim}}, \hat{P}(k')_{\text{sim}})$ be the measured covariance of signal plus noise simulations of the data pipeline. Put the square root of Gaussian errors derived for the simulation (Equation C5) along the diagonal of $\mathbf{\Lambda}_{\text{sim}}$ and likewise for the measured data $\mathbf{\Lambda}_{\text{data}}$. Then, recalibrate the band power covariance measured in simulations against the variance measured in the data as

$$\mathbf{C}_{\text{data}} = \mathbf{\Lambda}_{\text{data}} \mathbf{\Lambda}_{\text{sim}}^{-1} \mathbf{C}_{\text{sim}} \mathbf{\Lambda}_{\text{sim}}^{-1} \mathbf{\Lambda}_{\text{data}}. \quad (\text{C7})$$

In the operation $\mathbf{\Lambda}_{\text{sim}}^{-1} \mathbf{\Lambda}_{\text{data}}$, the common factor $\nu(k)$ drops out of the Gaussian errors of the data and simulations and does not need to be calculated explicitly. The simulation pipeline does not need a high-fidelity model of the real data's covariance. Instead, the goal of the simulations is to measure the off-diagonal terms, and the $A \times A$ combinations of the real data give Gaussian errors that calibrate the amplitudes.

REFERENCES

- Ali, Z. S., Parsons, A. R., Zheng, H., et al. 2015, ArXiv e-prints, arXiv:1502.06016
- Alonso, D., Bull, P., Ferreira, P. G., & Santos, M. G. 2015, MNRAS, 447, 400
- Bagla, J. S., & Loeb, A. 2009, ArXiv e-prints, arXiv:0905.1698
- Bandura, K., Addison, G. E., Amiri, M., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9145, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 22
- Battye, R. A., Brown, M. L., Browne, I. W. A., et al. 2012, ArXiv e-prints, arXiv:1209.1041
- Bebbington, D. H. O. 1986, MNRAS, 218, 577
- Bigot-Sazy, M.-A., Dickinson, C., Battye, R. A., et al. 2015, ArXiv e-prints, arXiv:1507.04561
- Blake, C., Brough, S., Colless, M., et al. 2011, MNRAS, 415, 2876
- Breyse, P. C., Kovetz, E. D., & Kamionkowski, M. 2014, MNRAS, 443, 3506
- . 2015, ArXiv e-prints, arXiv:1503.05202
- Bryerton, E. W. 2011, IEEE Transactions on Microwave Theory Techniques, 59, 2117
- Carilli, C. L. 2011, ApJ, 730, L30
- Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, Nature, 466, 463
- Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, Physical Review Letters, 100, 091303
- Crites, A. T., Bock, J. J., Bradford, C. M., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9153, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1
- Croft, R. A. C., Miralda-Escudé, J., Zheng, Z., et al. 2015, ArXiv e-prints, arXiv:1504.04088
- Das, S., Marriage, T. A., Ade, P. A. R., et al. 2011, ApJ, 729, 62
- de Putter, R., Holder, G. P., Chang, T.-C., & Dore, O. 2014, ArXiv e-prints, arXiv:1403.3727
- Dillon, J. S., Liu, A., & Tegmark, M. 2013, Phys. Rev. D, 87, 043005
- Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, Phys. Rev. D, 89, 023002
- Doré, O., Bock, J., Ashby, M., et al. 2014, ArXiv e-prints, arXiv:1412.4872
- Drinkwater, M. J., Jurek, R. J., Blake, C., et al. 2010, MNRAS, 401, 1429
- Dünner, R., Hasselfield, M., Marriage, T. A., et al. 2013, ApJ, 762, 10
- Eales, S., Dunne, L., Clements, D., et al. 2010, PASP, 122, 499
- Efstathiou, G., Gratton, S., & Paci, F. 2009, MNRAS, 397, 1355
- Gong, Y., Cooray, A., Silva, M., et al. 2012, ApJ, 745, 49
- Gong, Y., Silva, M., Cooray, A., & Santos, M. G. 2014, ApJ, 785, 72
- Hamilton, A. J. S., & Tegmark, M. 2000, MNRAS, 312, 285
- Harnois-Déraps, J., & Pen, U.-L. 2012, MNRAS, 423, 2288
- Hauser, M. G., & Dwek, E. 2001, ARA&A, 39, 249
- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, ApJ, 567, 2
- Hogan, C. J., & Rees, M. J. 1979, MNRAS, 188, 791
- Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, Phys. Rev. D, 67, 043004
- Hütsi, G., Gilfanov, M., & Sunyaev, R. 2012, A&A, 547, A21
- Kasai, M. 2010, ArXiv e-prints, arXiv:1012.2671
- Kaufman, J. P., Miller, N. J., Shimon, M., et al. 2014, Phys. Rev. D, 89, 062006
- Komatsu, E., Dunkley, J., Nolte, M. R., et al. 2009, ApJS, 180, 330
- Kusakabe, M., & Kawasaki, M. 2012, MNRAS, 419, 873
- Levenson, L. R., & Wright, E. L. 2008, ApJ, 683, 585
- Li, T. Y., Wechsler, R. H., Devaraj, K., & Church, S. E. 2015, ArXiv e-prints, arXiv:1503.08833
- Lidz, A., Furlanetto, S. R., Oh, S. P., et al. 2011, ApJ, 741, 70
- Liu, A., Parsons, A. R., & Trott, C. M. 2014, Phys. Rev. D, 90, 023019
- Liu, A., & Tegmark, M. 2011, Phys. Rev. D, 83, 103006
- . 2012, MNRAS, 419, 3491
- Loeb, A., & Wyithe, J. S. B. 2008, Physical Review Letters, 100, 161301
- Madau, P., Meiksin, A., & Rees, M. J. 1997, ApJ, 475, 429
- Masui, K. W. 2013, PhD thesis, University of Toronto (Canada)
- Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, ApJ, 763, L20
- McQuinn, M., & Switzer, E. R. 2009, Phys. Rev. D, 80, 063010
- Moore, D. F., Aguirre, J. E., Parsons, A. R., Jacobs, D. C., & Pober, J. C. 2013, ApJ, 769, 154
- Murthy, J., Henry, R. C., & Sujatha, N. V. 2010, ApJ, 724, 1389
- Nityananda, R. 2010, NCRA Technical Reports
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, MNRAS, 433, 639
- Pearson, T. J., Mason, B. S., Readhead, A. C. S., et al. 2003, ApJ, 591, 556
- Pober, J. C., Parsons, A. R., DeBoer, D. R., et al. 2013, AJ, 145, 65
- Pullen, A. R., Chang, T.-C., Doré, O., & Lidz, A. 2013, ApJ, 768, 15
- Pullen, A. R., Doré, O., & Bock, J. 2014, ApJ, 786, 111
- Righi, M., Hernández-Monteagudo, C., & Sunyaev, R. A. 2008, A&A, 489, 489
- Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, ApJ, 743, 90
- Scott, D., & Rees, M. J. 1990, MNRAS, 247, 510
- Seljak, U. 1998, ApJ, 503, 492

- Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014a, *ApJ*, 781, 57
- Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2014b, *ArXiv e-prints*, arXiv:1401.2095
- Sigurdson, K., & Furlanetto, S. R. 2006, *Physical Review Letters*, 97, 091301
- Silva, M. B., Santos, M. G., Cooray, A., & Gong, Y. 2014, *ArXiv e-prints*, arXiv:1410.4808
- Smith, R. E., Peacock, J. A., Jenkins, A., et al. 2003, *MNRAS*, 341, 1311
- Switzer, E. R., & Liu, A. 2014, *ApJ*, 793, 102
- Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *MNRAS*, 434, L46
- Takahashi, Y. D., Barkats, D., Battle, J. O., et al. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7020, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1
- Tegmark, M. 1997, *Phys. Rev. D*, 55, 5895
- Tegmark, M., Hamilton, A. J. S., Strauss, M. A., Vogeley, M. S., & Szalay, A. S. 1998, *ApJ*, 499, 555
- Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, *ApJ*, 480, 22
- Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005, *MNRAS*, 358, 833
- Uzgil, B. D., Aguirre, J. E., Bradford, C. M., & Lidz, A. 2014, *ApJ*, 793, 116
- Visbal, E., & Loeb, A. 2010, *JCAP*, 11, 16
- Wang, X., Tegmark, M., Santos, M. G., & Knox, L. 2006, *ApJ*, 650, 529
- Wolz, L., Abdalla, F. B., Blake, C., et al. 2014, *MNRAS*, 441, 3271
- Yue, B., Ferrara, A., Pallottini, A., Gallerani, S., & Vallini, L. 2015, *ArXiv e-prints*, arXiv:1504.06530